ANSYS CONVERGENCE CONFERENCES

2018年7月11-13日 上海

HPC高性能计算最佳实践

ANSYS Inc.

内容提要

- HPC 简介
- ANSYS HPC 软件配置
- ANSYS HPC 硬件选择
- 应用案例

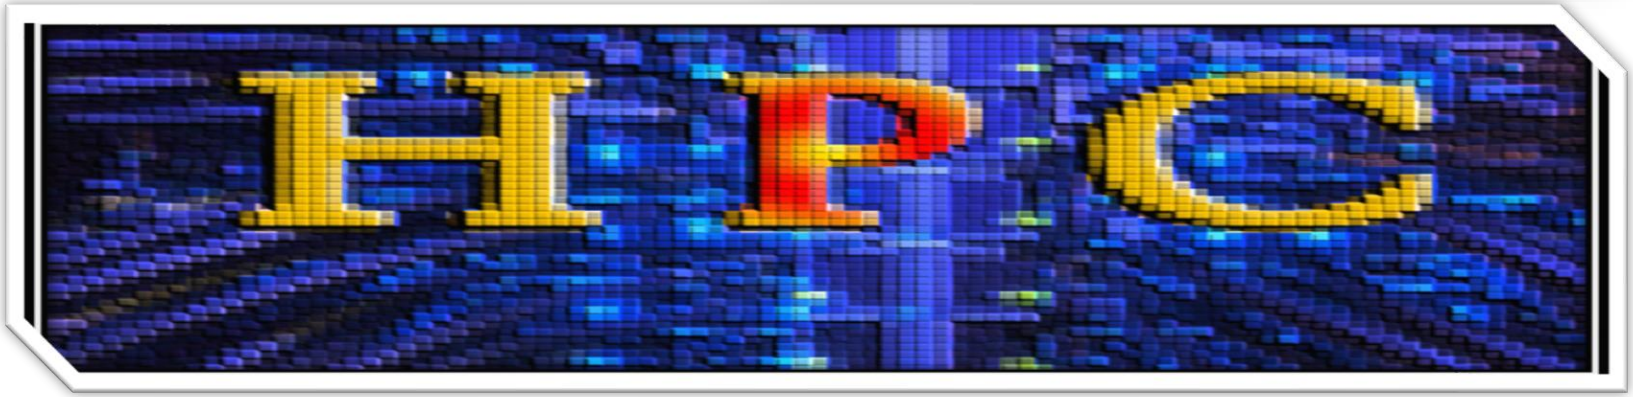

HPC简介

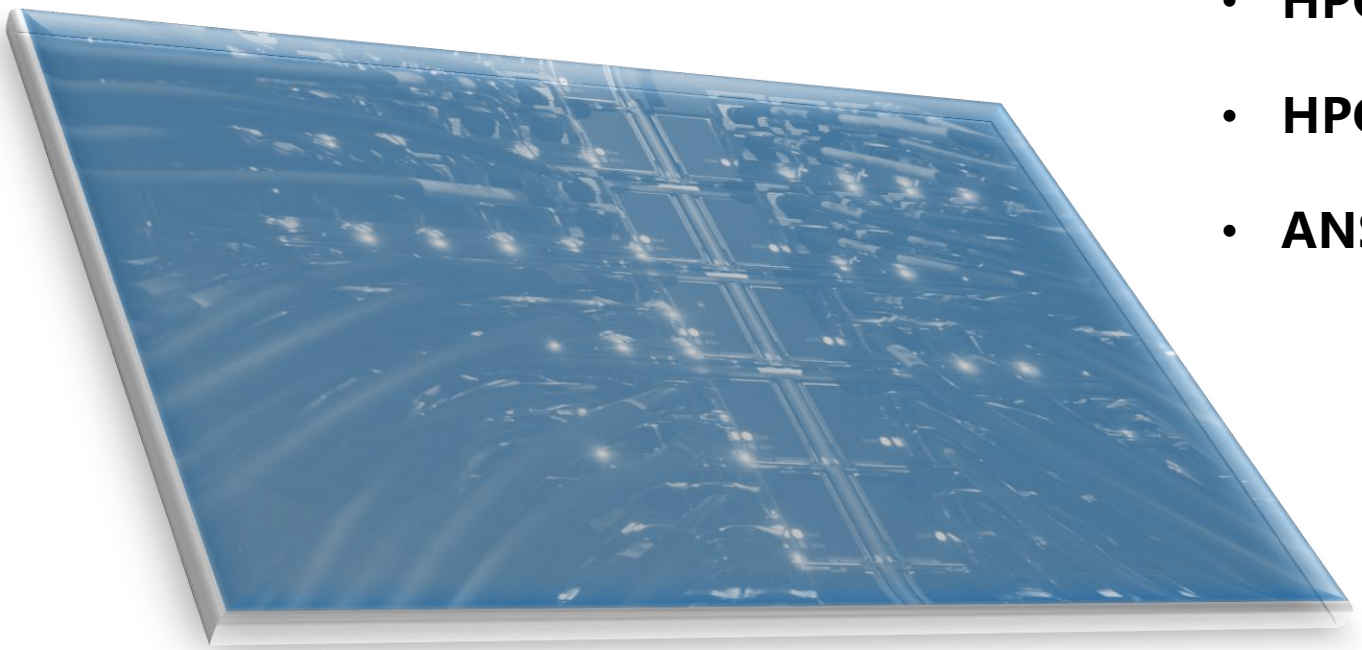

- 什么是HPC
- HPC带来的好处
- HPC计算原理
- ANSYS HPC加速效果

什么是高性能计算（HPC）

- *高性能计算一般是指通过集合更多的计算资源提供远远超过单一工作站的计算能力去求解科学、工程问题的实践。*

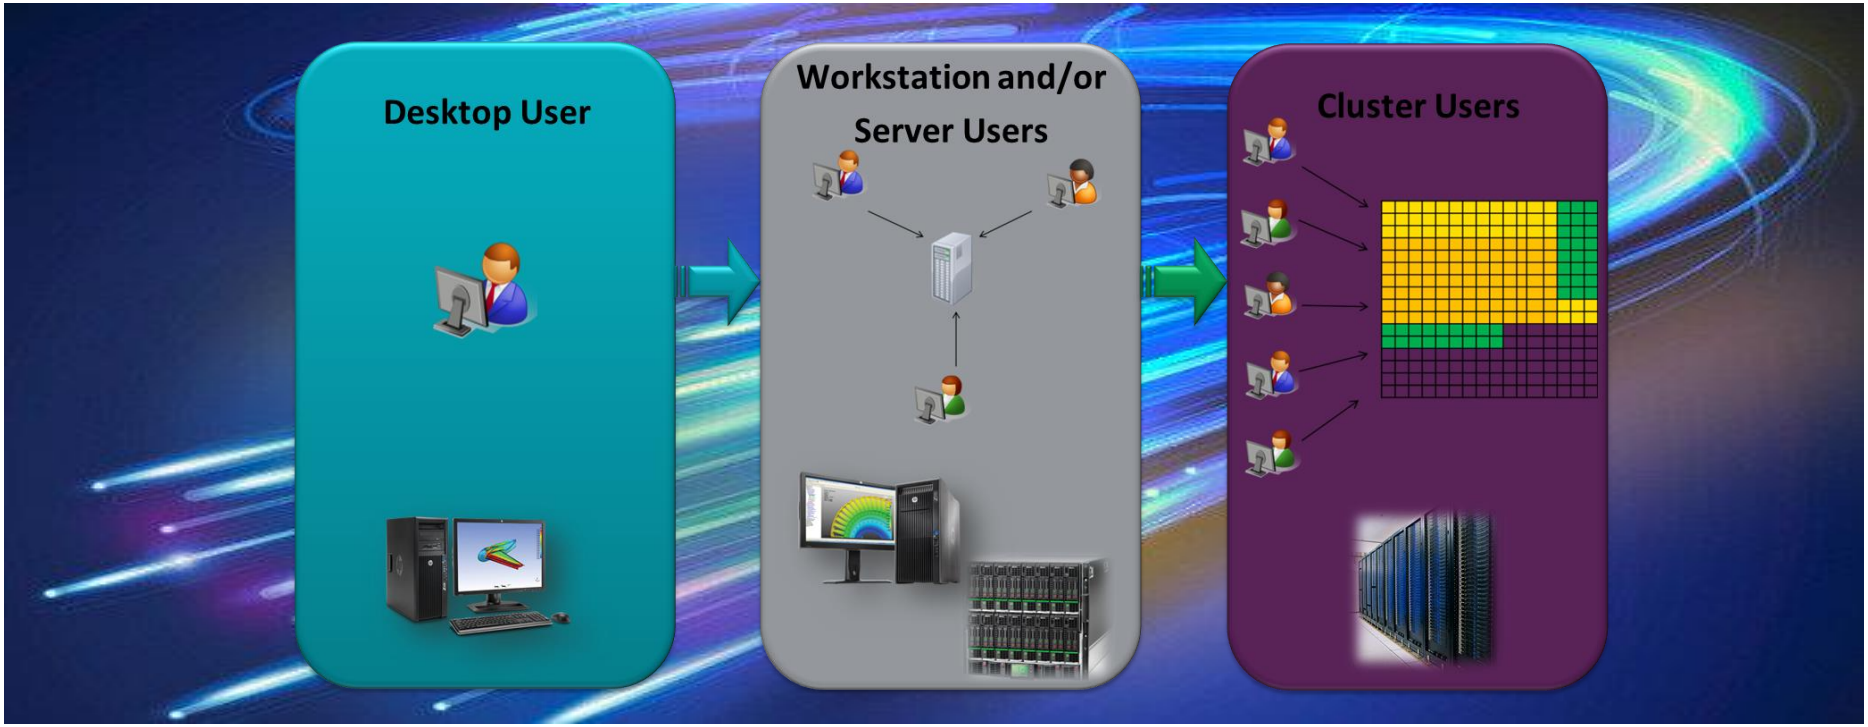

HPC带来的好处

提高保真度

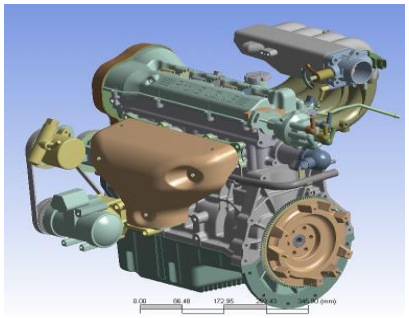

计算更复杂的装配体

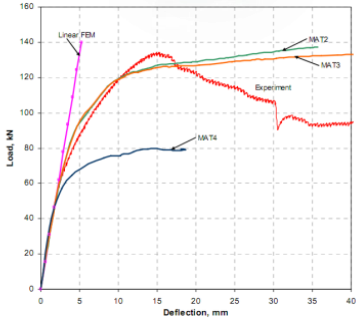

考虑更多的非线性

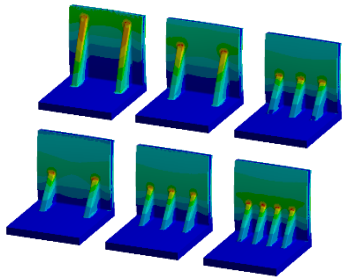

更多的设计场景验证

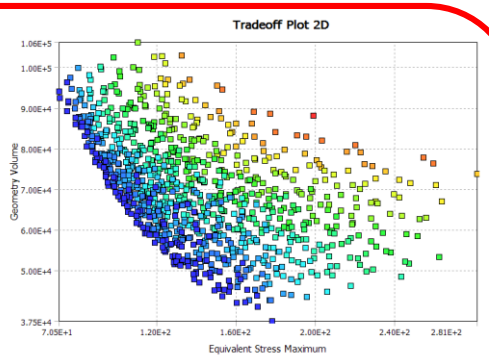

更多的优化分析

HPC计算原理

ANSYS HPC:

通过将大规模计算问题分解成可以并行计算的子问题，分配到多个计算核心（CPU或者GPU）上进行并行计算。充分利用计算资源，加速计算。

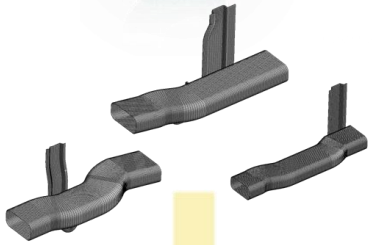

ANSYS HPC Parametric Pack：

将参数化模型的每一组参数设置分配到多个计算节点，同时求解多个不同参数设置的模型，实现加速。可以结合HPC一同使用。

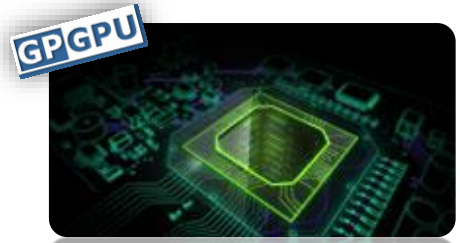

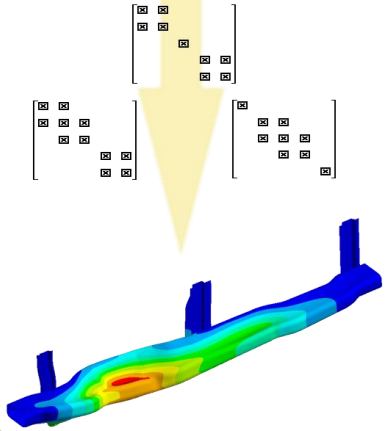

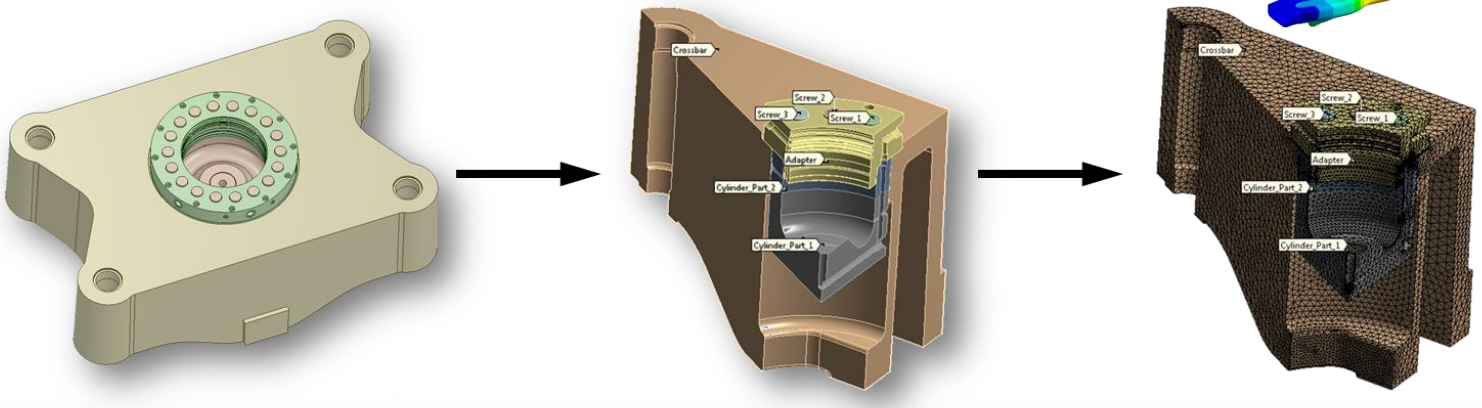

加速效果-HPC

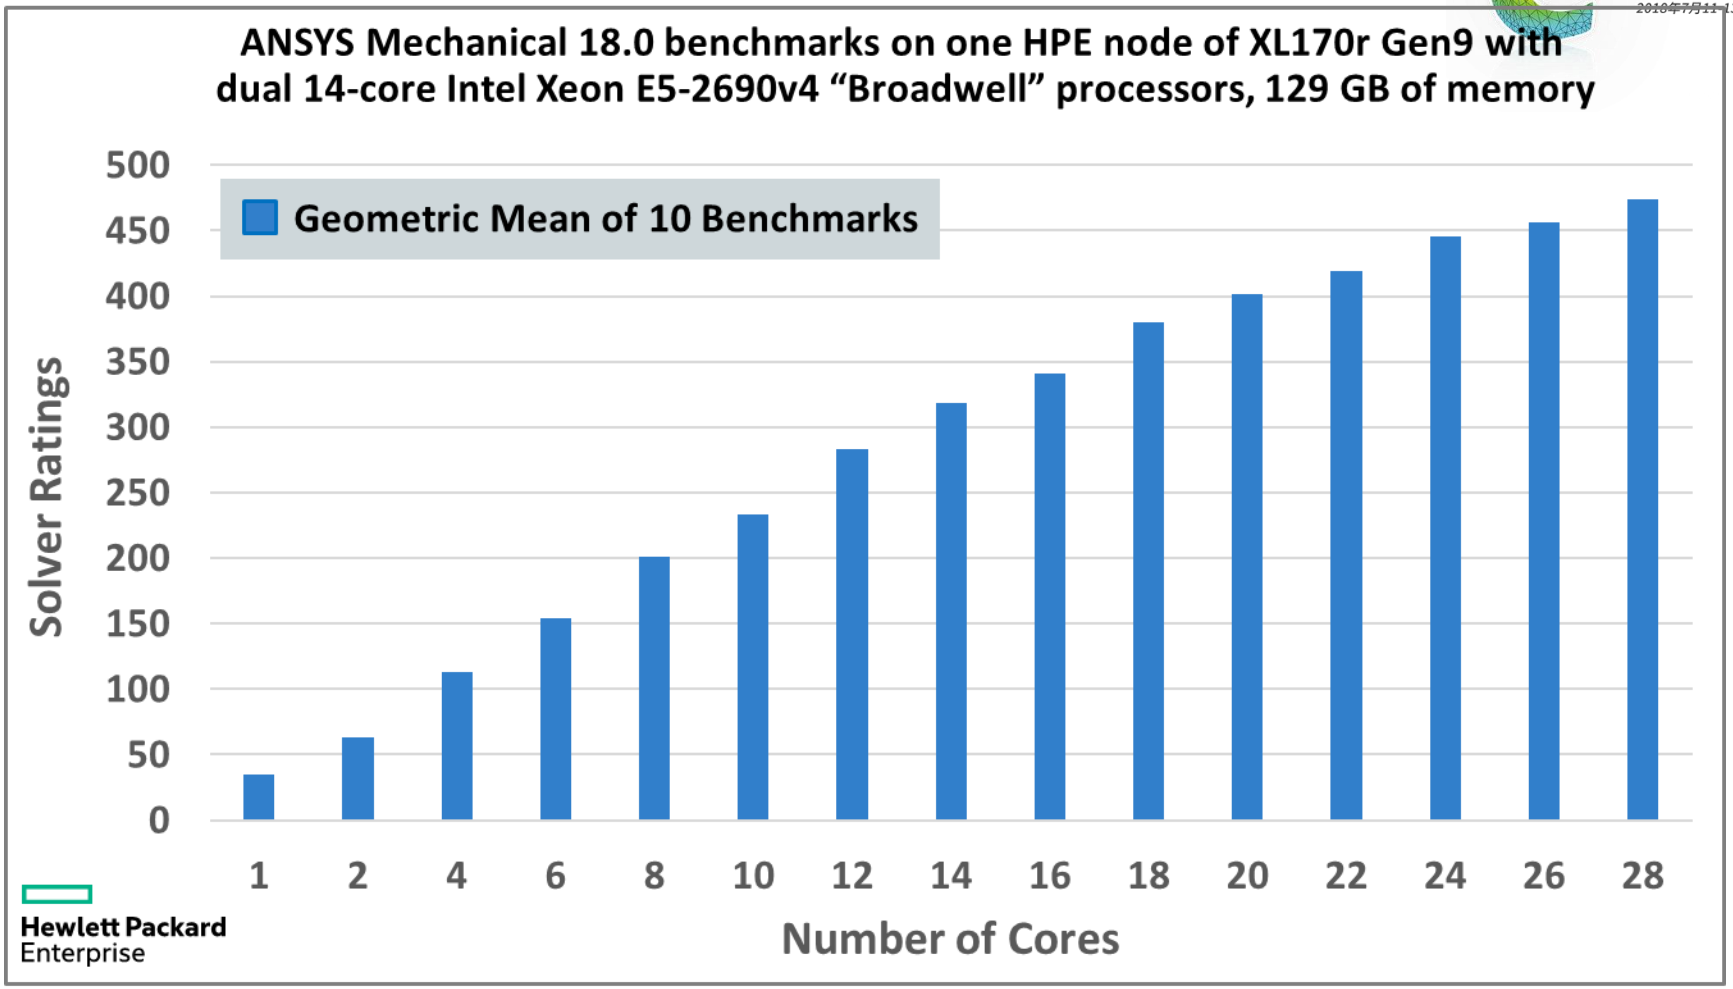

Sovler Rating：It is computed by dividing the number of seconds in a day (86400 seconds) by the number of seconds required to run the benchmark. A higher rating means faster performance.

加速效果-HPC Parametric Pack

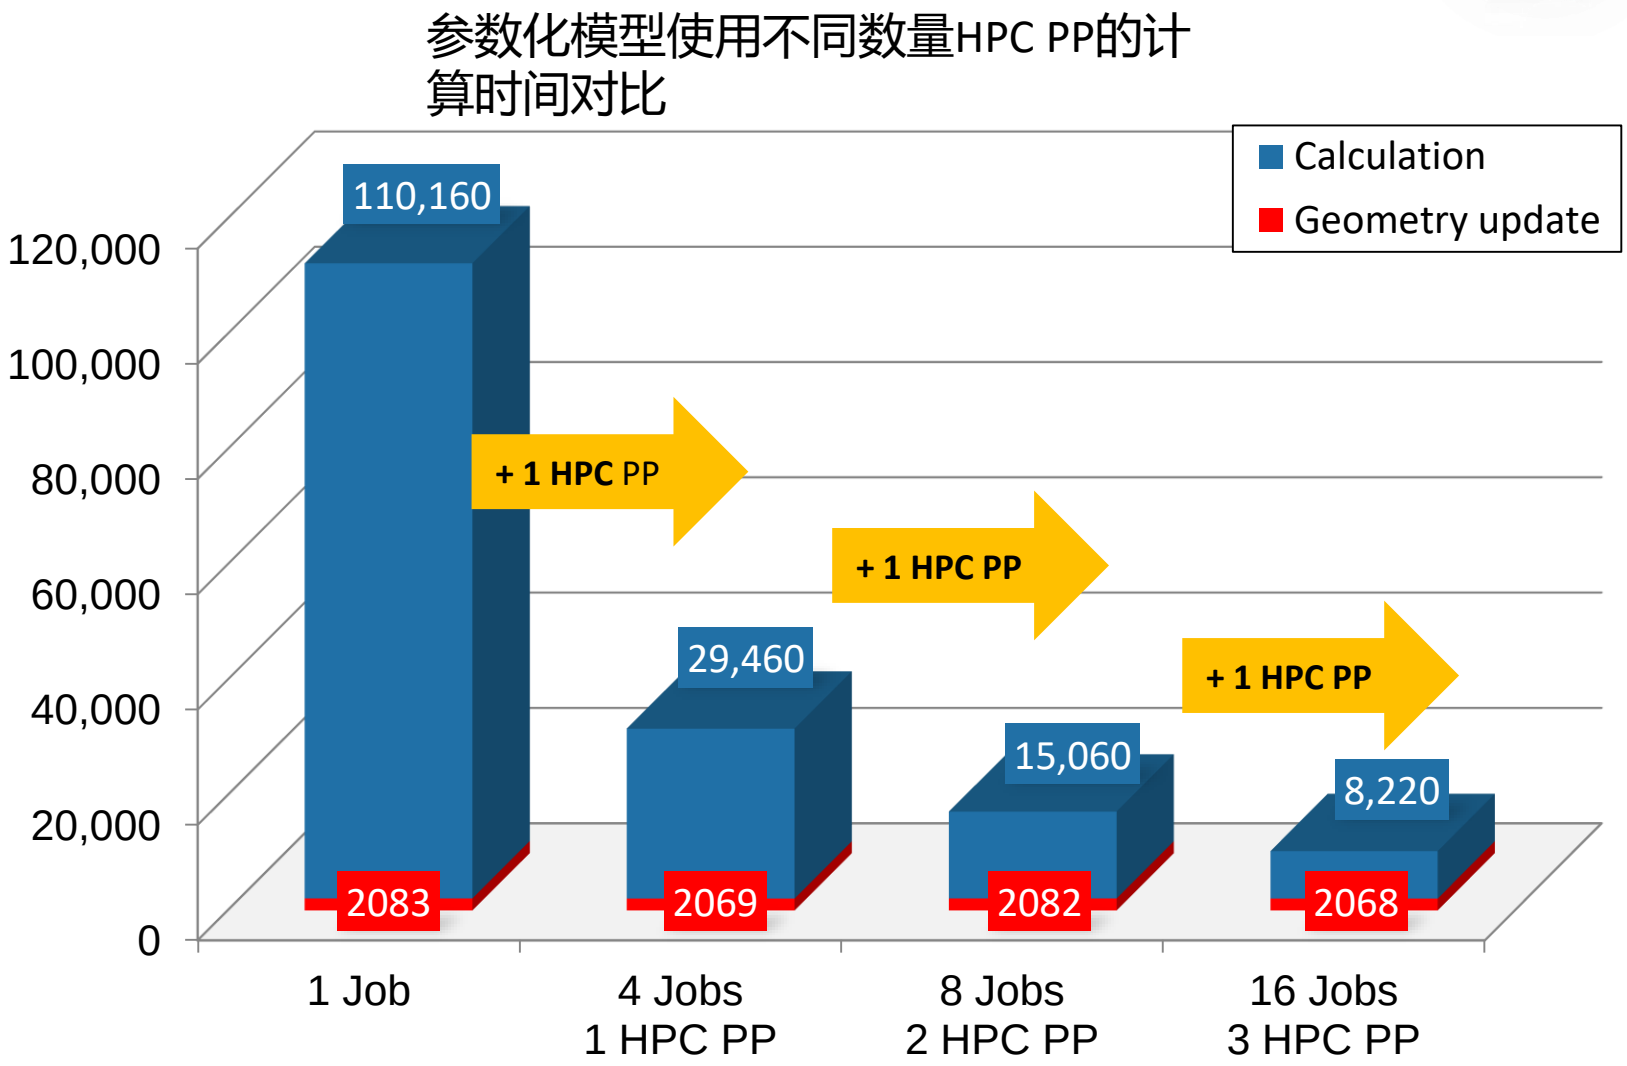

HPC软件配置

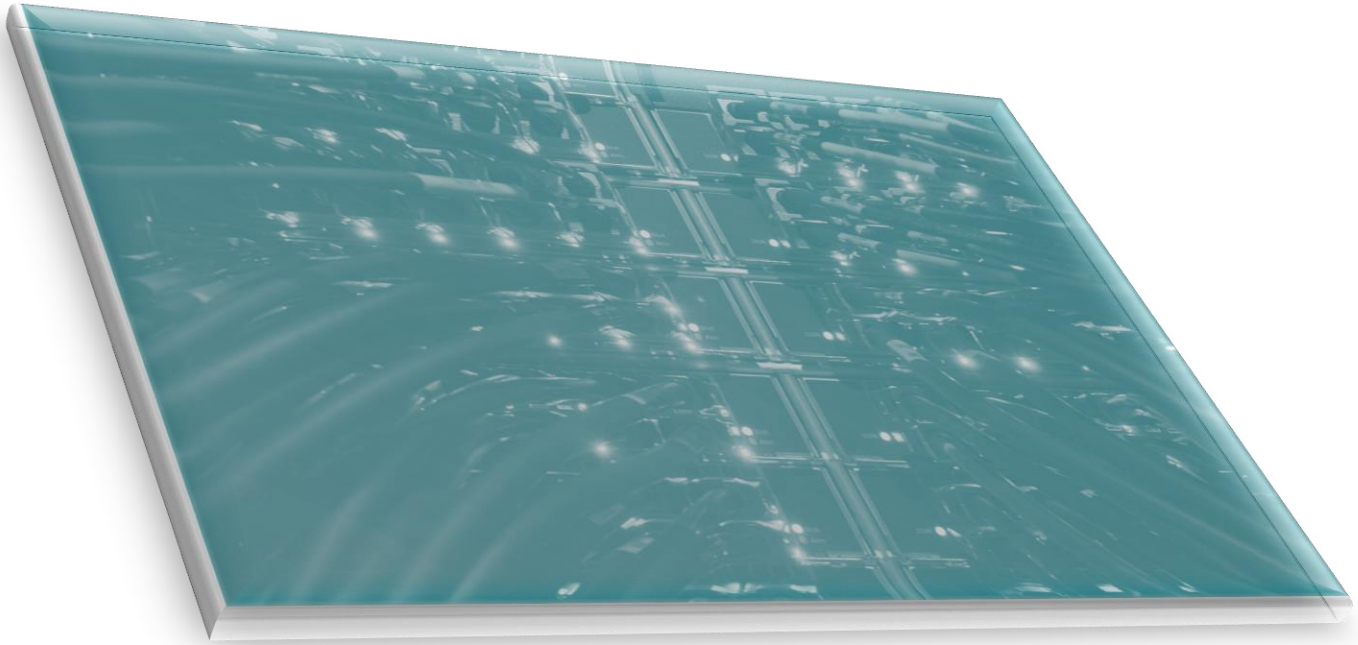

软件配置选项

HPC (per-process)

HPC Pack

- HPC product rewarding volume parallel processing for high-fidelity simulations
- Each simulation consumes one or more Packs
- Parallel enabled increases quickly with added Packs

HPC Workgroup

- HPC product rewards volume parallel processing for increased simulation throughput shared among engineers throughout a single location or the world
- 16 to 32768 parallel shared across any number of simulations on a single server

HPC Parametric Pack

- Enables simultaneous execution of multiple design points while consuming just one set of licenses

Single HPC solution for FEA/CFD/FSI and any level of fidelity

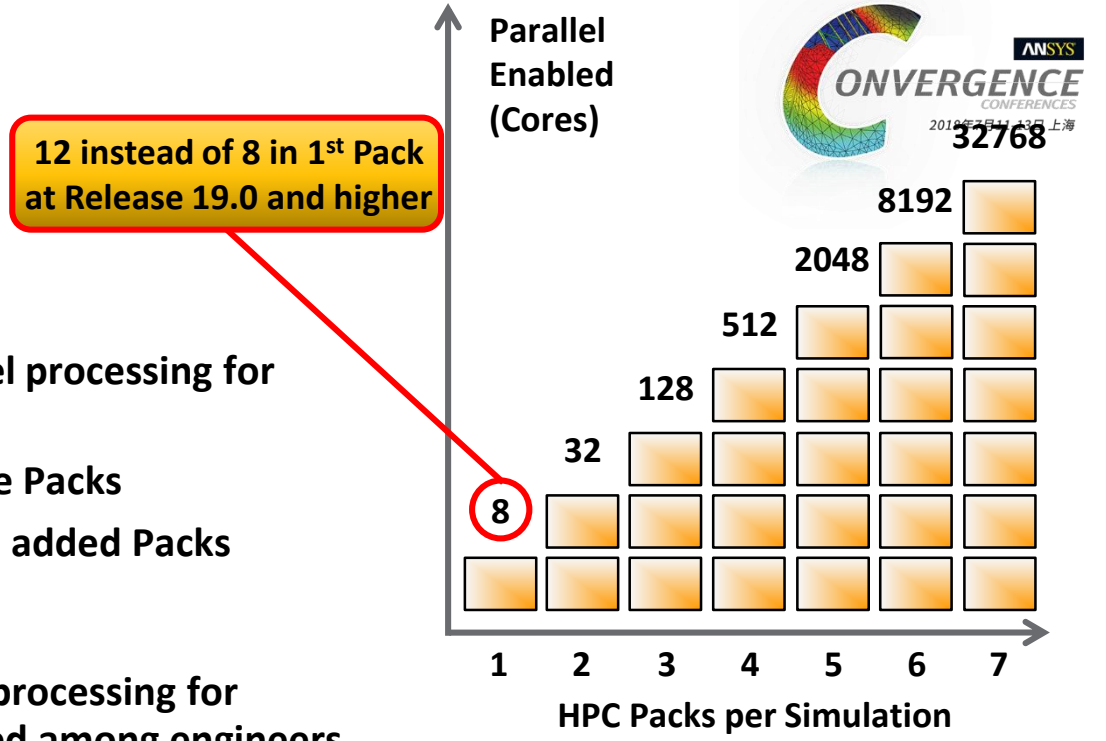

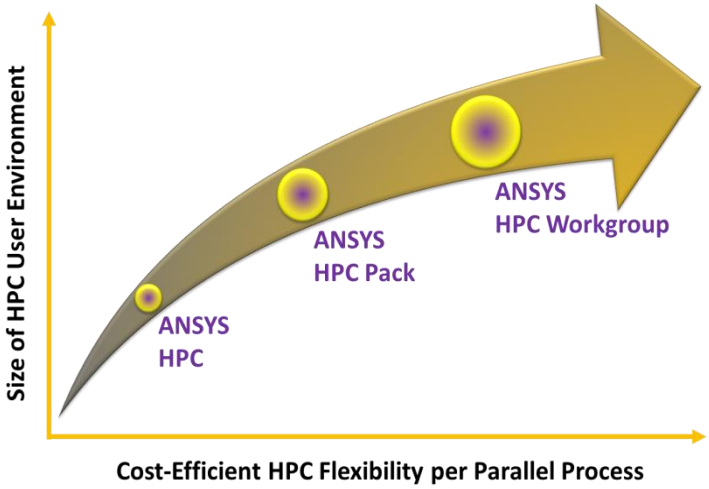

ANSYS 19.0新特性

More products are now using ANSYS HPC

- Standalone HPC licenses, HPC Packs and HPC Workgroup become **more flexible** and **work across physics** with all ANSYS Mechanical, Fluids and Electronics products*

Note: R19.0 license manager is required. For ANSYS Mechanical and Fluids products changes are backward compatible; for ANSYS Electronics products changes are compatible with version 19.0 and forward

4 Built-in HPCs now across all physics

- **4 built-in HPCs** are now included in Mechanical, Fluids and Electronics products, including ANSYS AIM and ANSYS Chemkin Enterprise.

Note: built-in HPCs are linked to a solver seat and cannot be shared with other solver seats!

HPC Packs are now additive

- HPC Packs becomes **additive in nature** to the 4 built-in HPCs (e.g. 1 HPC Pack licenses 8 + 4 = 12 total cores, 2 HPC Pack license 32 + 4 = 36 total cores, etc.)

Note: the single, standalone HPCs are not additive to the Packs

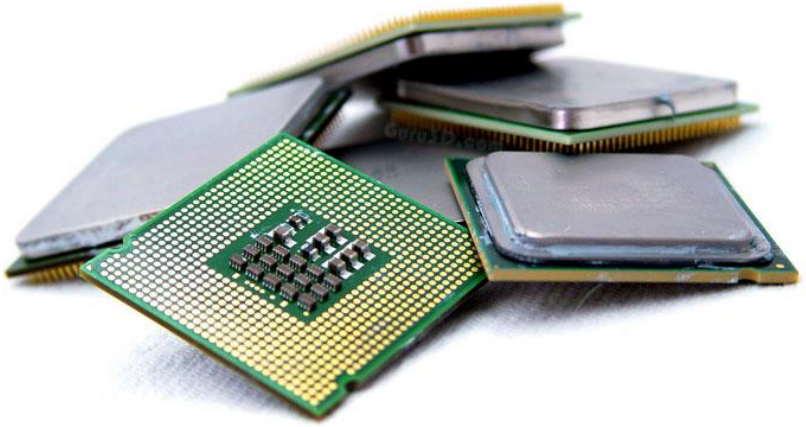

* Impacted products :

ANSYS Mechanical Pro, Premium, Enterprise	ANSYS CFD Premium and Enterprise
ANSYS Mechanical CFD	ANSYS HFSS
ANSYS AIM	ANSYS Q3D Extractor
ANSYS Maxwell	ANSYS Icepak
ANSYS Mechanical CFD Maxwell 3D	ANSYS Chemkin-Pro and Enterprise
ANSYS Mechanical Maxwell 3D	ANSYS SIwave

ANSYS HPC *Parametric* Pack介绍

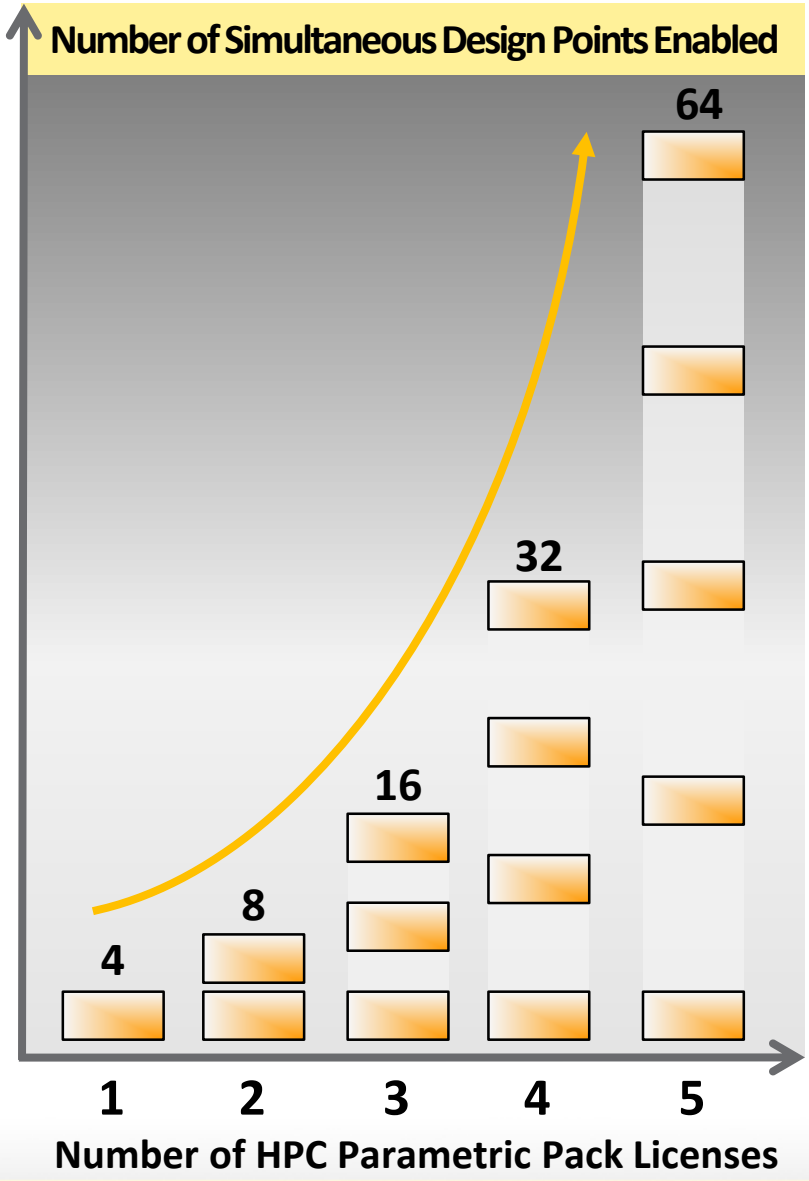

HPC license for running parametric FEA or CFD simulations on multiple CPU cores simultaneously, and more cost effectively

Key Benefits

- **Ability to automatically and simultaneously execute design points while consuming just one set of application licenses**
- **Scalable because number of simultaneous design points enabled increases quickly with added packs**
- **Amplifies complete workflow because design points can include execution of multiple applications (pre, meshing, solve, HPC, post)**

HPC Parametric Packs大幅缩短设计时间

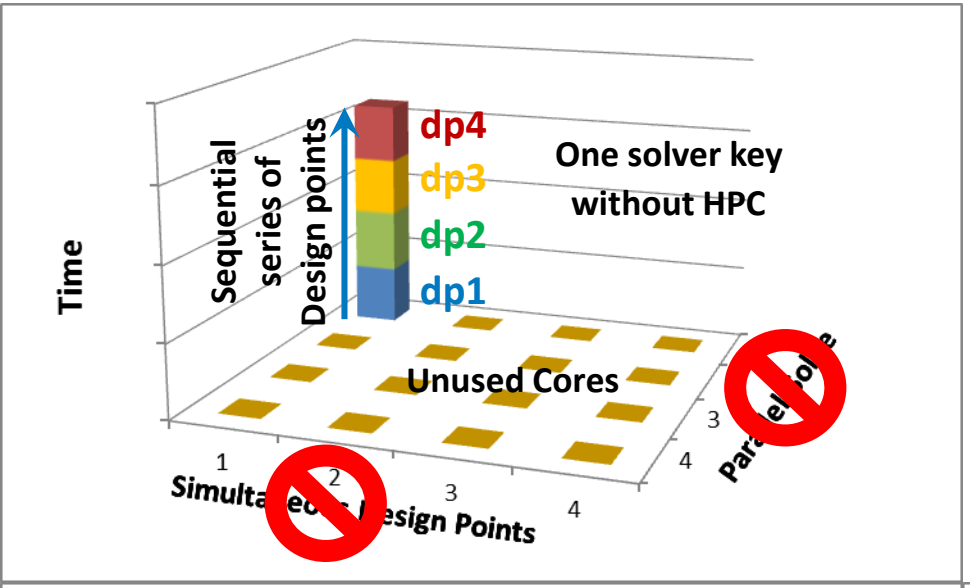

HPC Parametric Packs amplify both solver licenses and HPC licenses allowing you to drastically reduce time to innovation, without the cost of additional solver or HPC licenses…

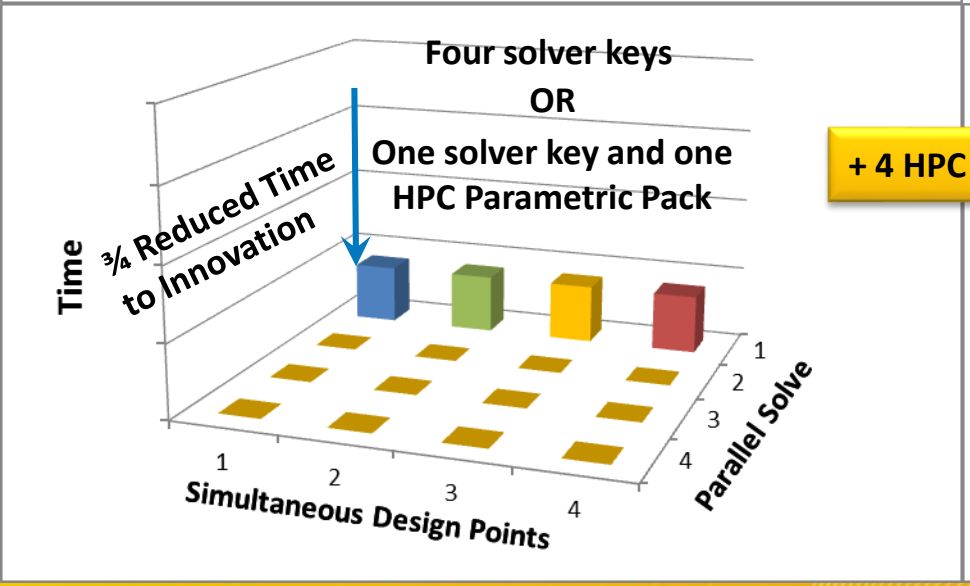

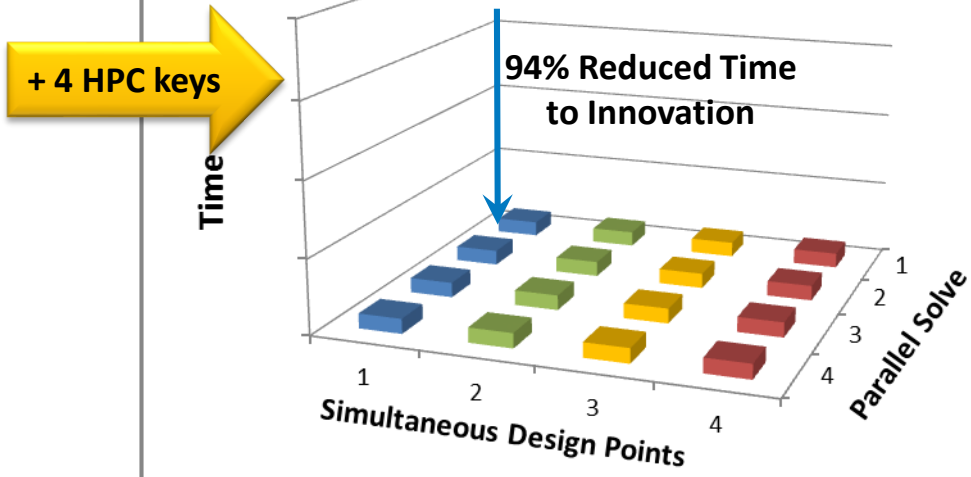

GPU加速

GPU acceleration can be enabled through all ANSYS HPC product licenses: ANSYS HPC, ANSYS HPC Pack and ANSYS HPC Workgroup.

Fluids / Structural products

1 GPU requires 1 HPC task as long as **GPUs ≤ CPU cores**

Examples:

- 2 HPC licenses enable up to 3 CPU cores + **3 GPUs** through the available 6 HPC tasks
- 1 HPC Pack enables up to 6 CPU cores + **6 GPUs** through the available 12 HPC tasks
- 2 HPC Packs enable up to 18 CPU cores + **18 GPUs** through the available 36 HPC tasks

Electronics products

1 GPU unlocked by every 8 HPC tasks

- 4 HPC licenses enable **1 GPU** through the available 8 HPC tasks
- 1 HPC Pack enables up to 12 CPU cores + **1 GPUs** through the available 12 HPC tasks
- 2 HPC Packs enable up to 36 CPU cores + **4 GPUs** through the available 36 HPC tasks

我该选择哪种配置？

- HPC license cost decreases as more are purchased either as HPC Packs or as HPC Workgroups.
- ANSYS HPC and ANSYS HPC Workgroup gives flexible use of a pool of licenses.
- ANSYS HPC Pack gives “quick” scale-up but is more restrictive in how users can use it.
- The ability to be more flexible is why the HPC Workgroup options cost more than the HPC Packs.
- HPC Parametric Pack enables more cost-effective licensing for design exploration and optimization.

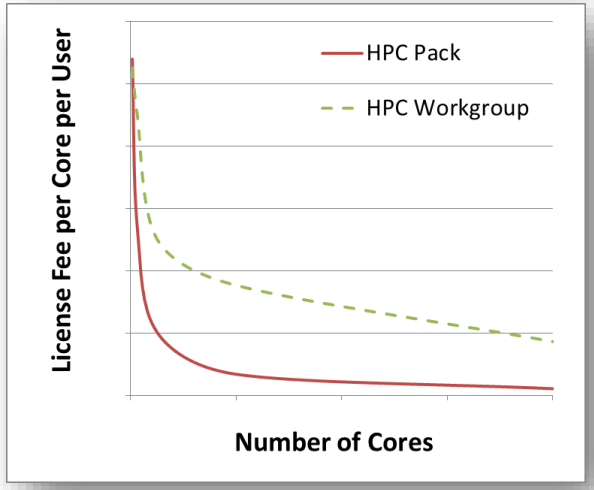

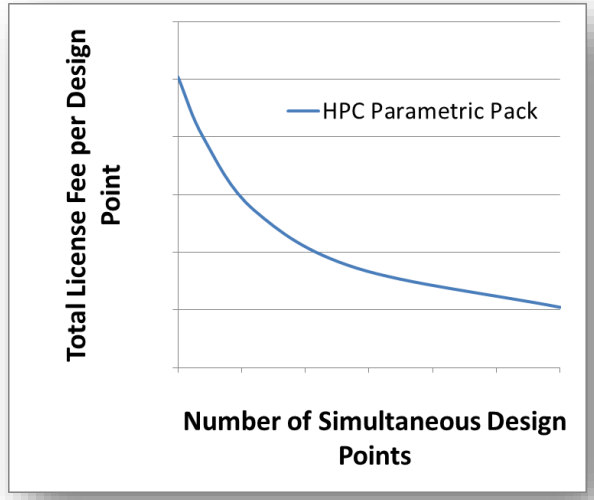

小结- 软件配置

Multiple licensing options to fit different requirements.	HPC Packs for quick scale-up.
	HPC Workgroup for Flexibility.
GPU's treated the same as cores in the licensing model.	
As you scale-up license cost decreases per core.	Per core pricing becomes less of an issue.

Running on 2,000 cores instead of 20 cores at 1.5X – and not 100X

Filling up a 1024- instead of 128-core cluster with 32-core jobs will cut the price per job in half!

Enabling 64 instead of 4 simultaneous design points at ~3X – and not 16X

ANSYS

选择什么样的硬件配置

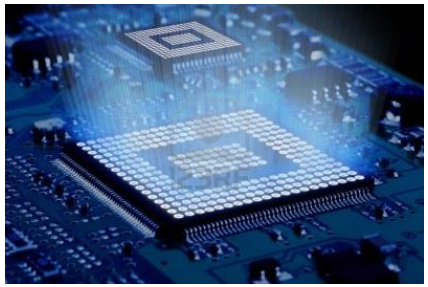

CPUs?

GPUs?

HDD vs. SSD

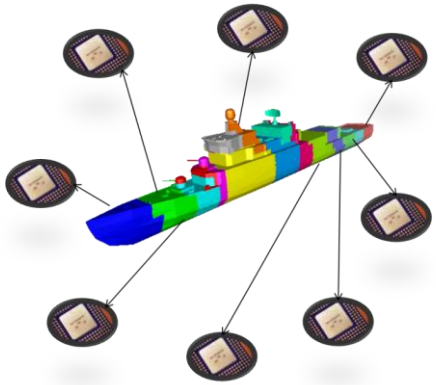

SMP vs. DMP

Clusters?

Interconnects?

HPC硬件术语

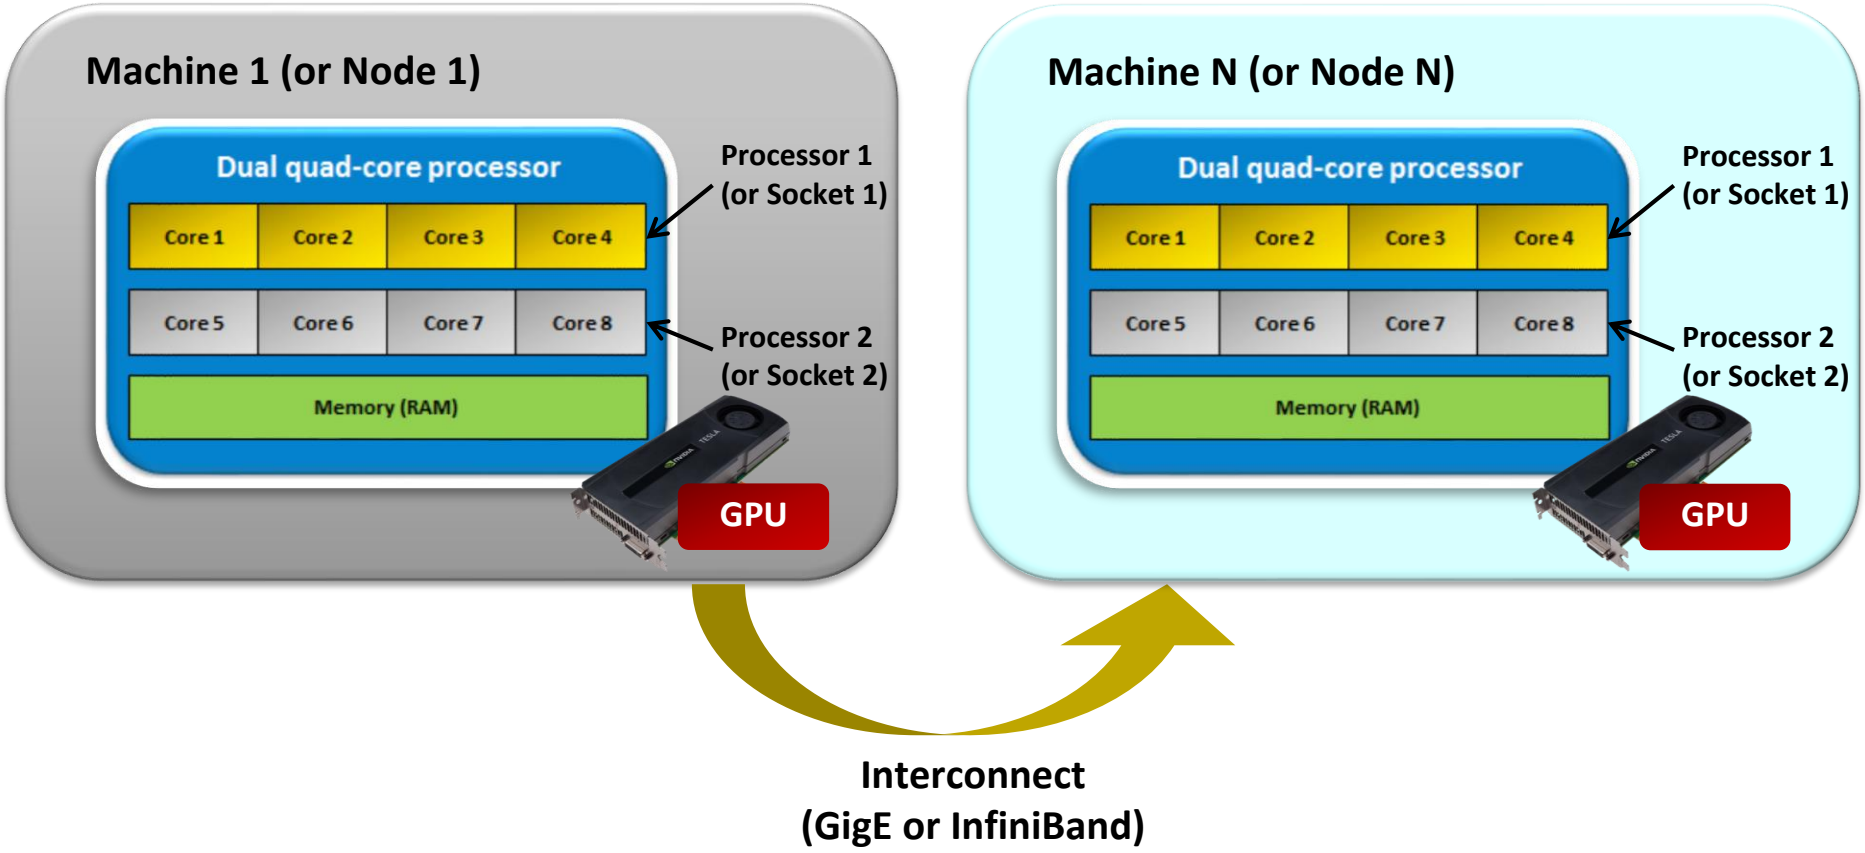
Machine 1 (or Node 1)
Dual quad-core processor
Core 1
Core 2
Core 3
Core 4
Core 5
Core 6
Core 7
Core 8
Memory (RAM)
Processor 1
(or Socket 1)
Processor 2
(or Socket 2)
GPU
Machine N (or Node N)
Dual quad-core processor
Core 1
Core 2
Core 3
Core 4
Core 5
Core 6
Core 7
Core 8
Memory (RAM)
Processor 1
(or Socket 1)
Processor 2
(or Socket 2)
GPU
Interconnect
(GigE or InfiniBand)

共享内存并行

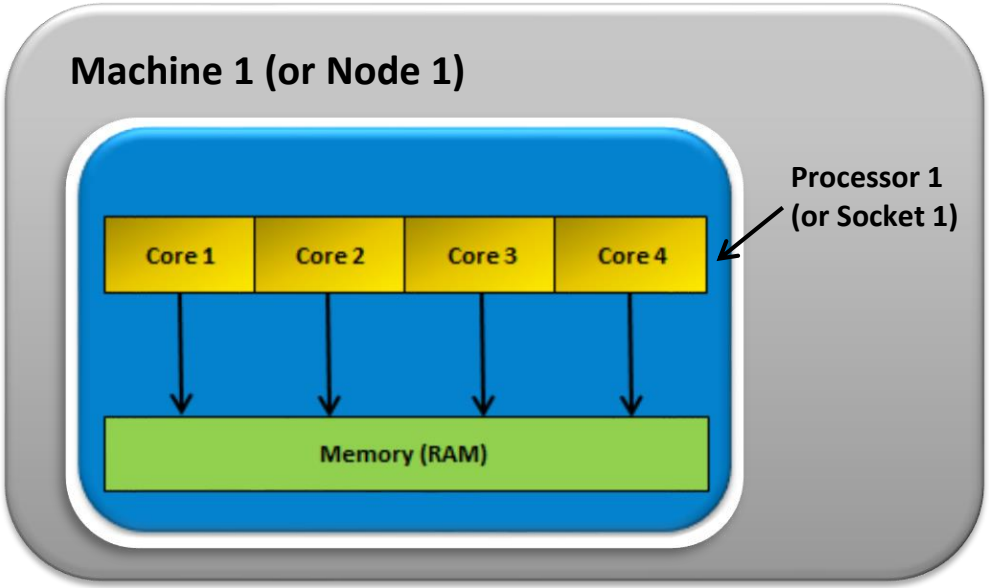

- **Single Machine Parallel (SMP) systems share a single global memory image that may be distributed physically across multiple cores, but is globally addressable.**
- **OpenMP is the industry standard.**

分布式内存并行

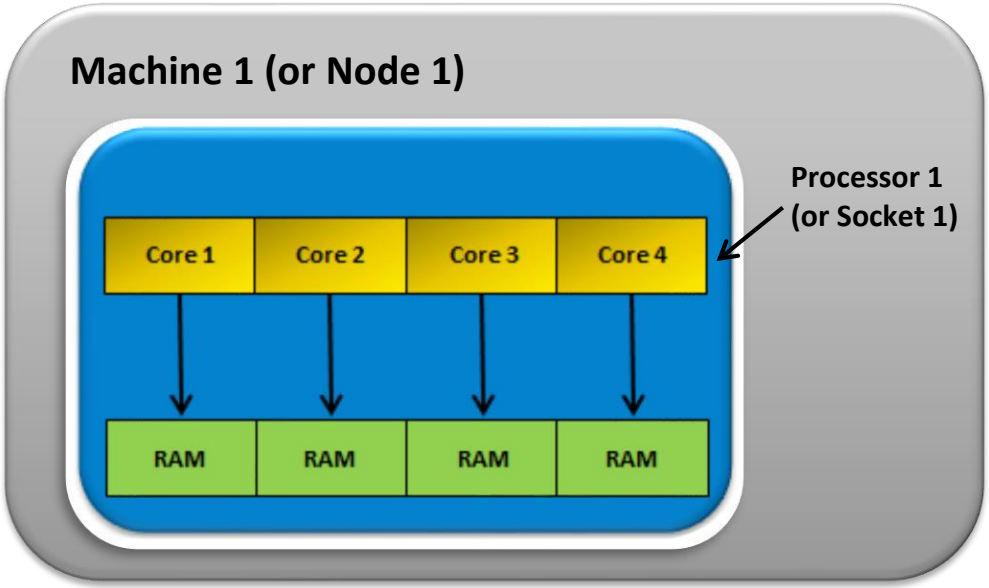

- Distributed memory parallel processing (DMP) assumes that physical memory for each process is separate from all other processes.
- Parallel processing on such a system requires some form of message passing software to exchange data between the cores.
- MPI (Message Passing Interface) is the industry standard for this.

了解时钟速度的影响

- ANSYS Mechanical

Using higher clock speed is always helpful to realize productivity gains

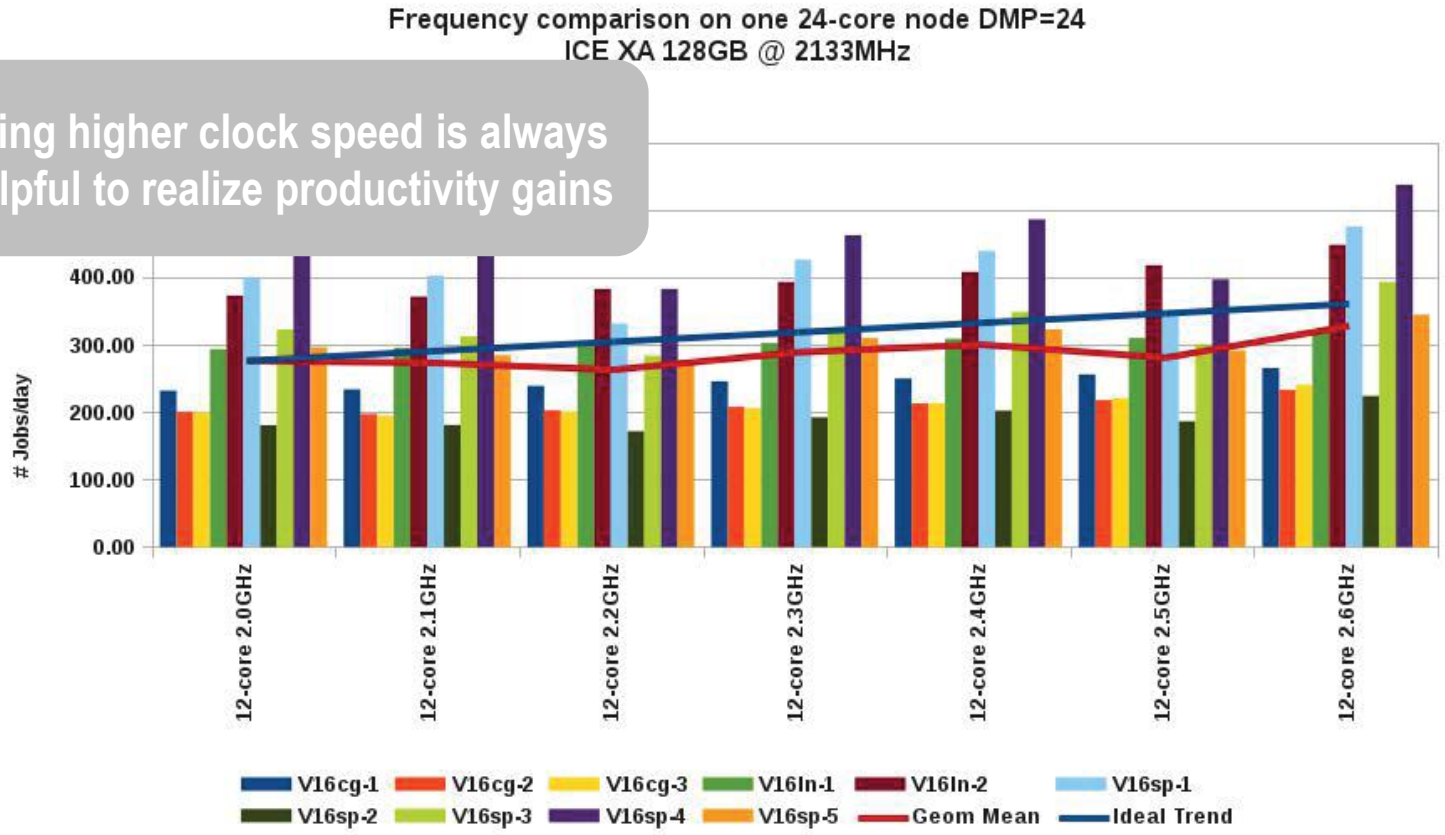

- Effect of increased core operating frequencies on the DMP benchmarks running on 12 cores
- Influence is highest for sparse solver benchmarks

了解内存带宽的影响

- Is 24 Cores Equal to 24 Cores?

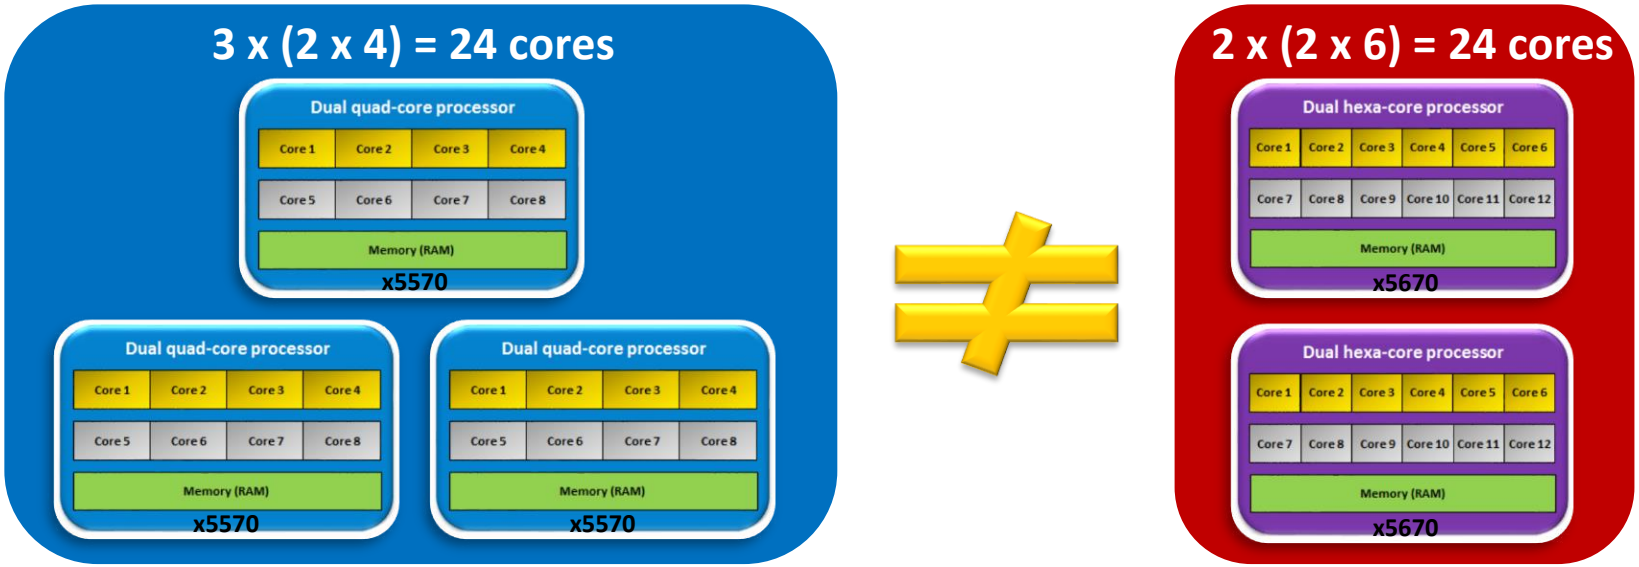

了解内存带宽的影响

- Is 24 Cores Equal to 24 Cores?

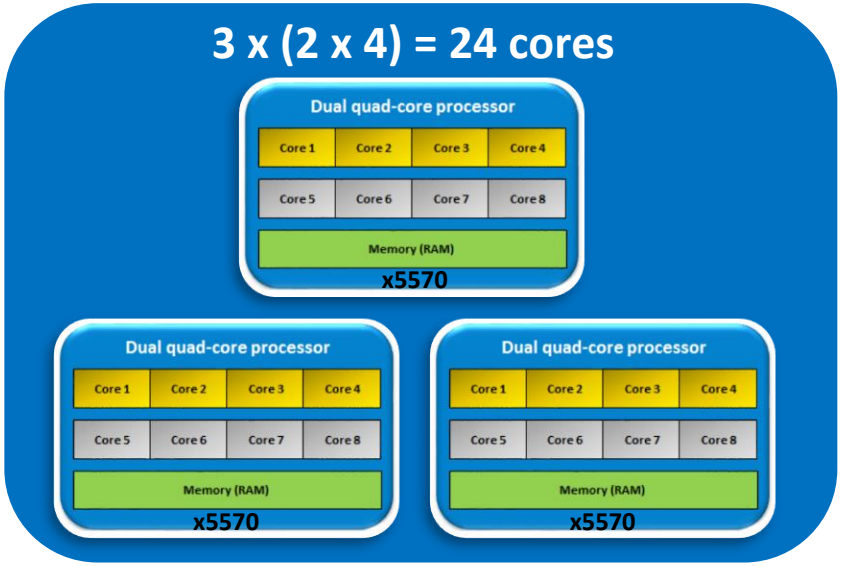

≠

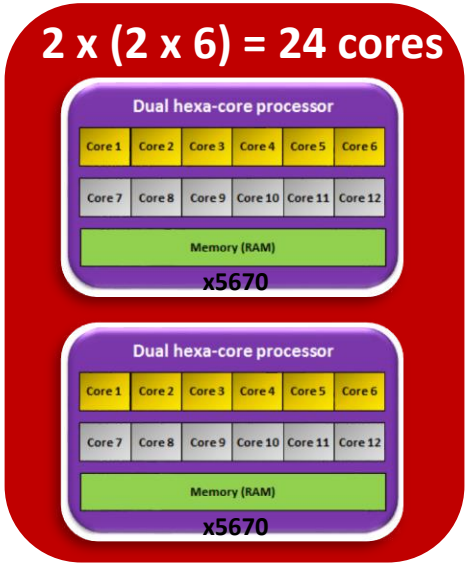

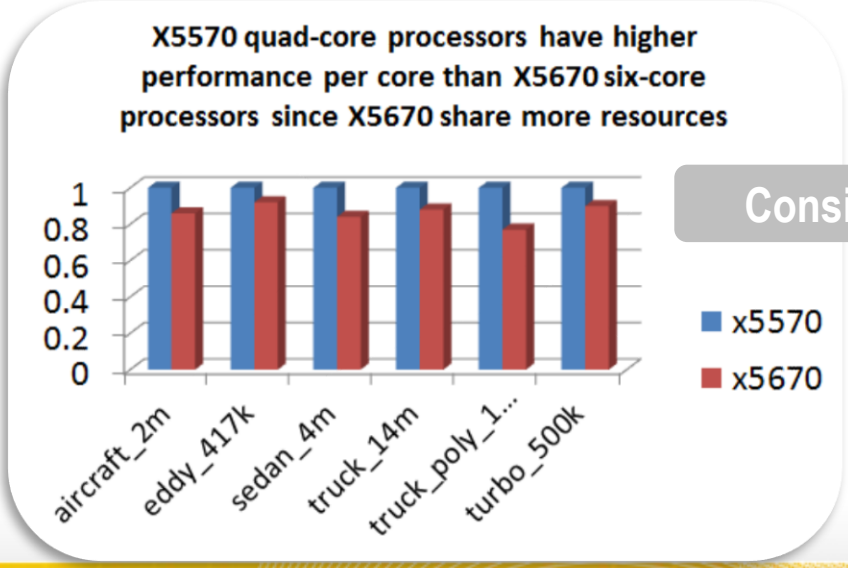

Consider memory per core!

分布式内存并行优于共享内存并行

SMP vs. DMP

Speedup Factor vs. Number of Cores for ANSYS Mechanical

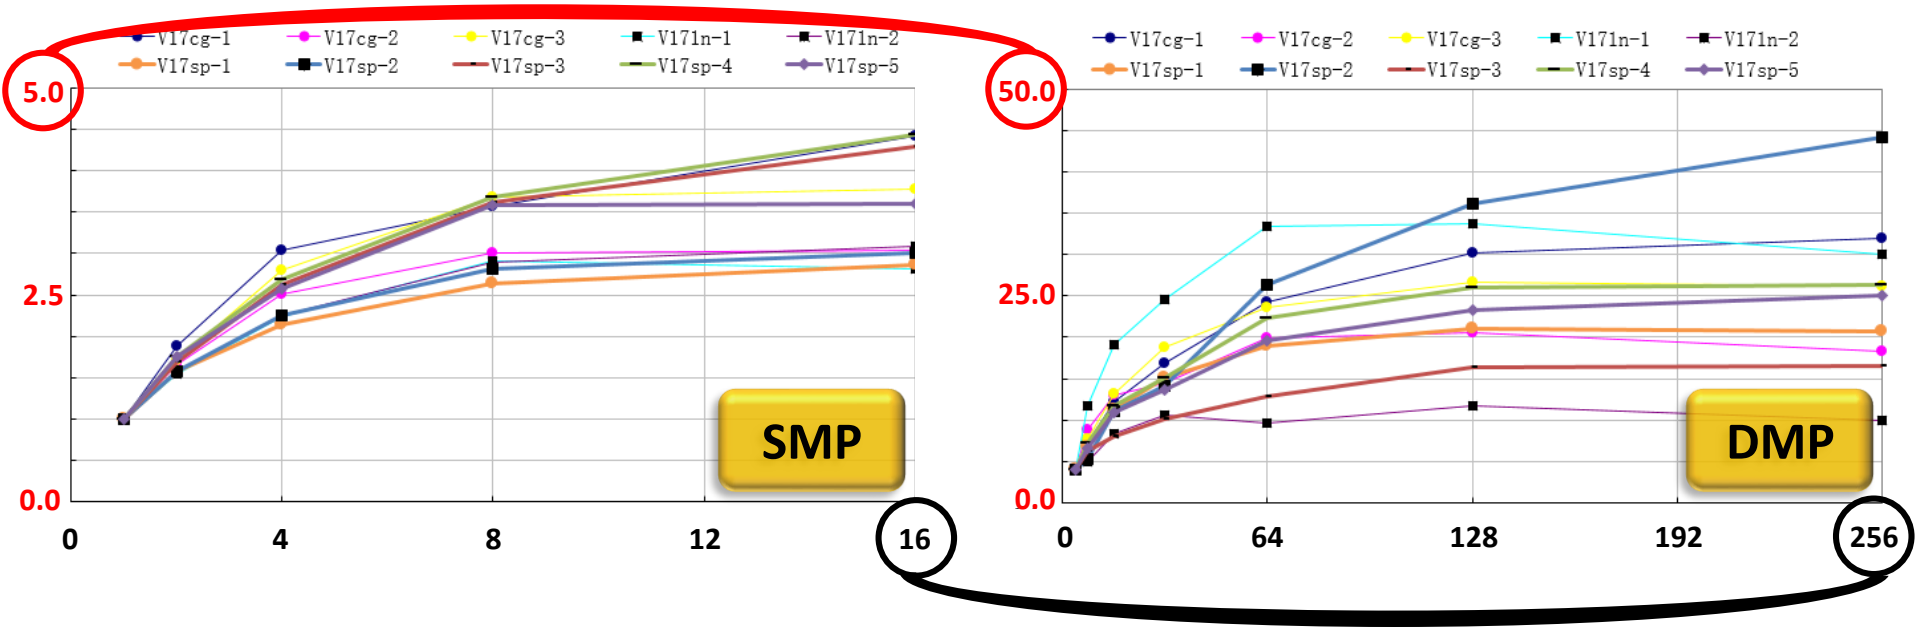

GPU 加速

ANSYS Application Examples

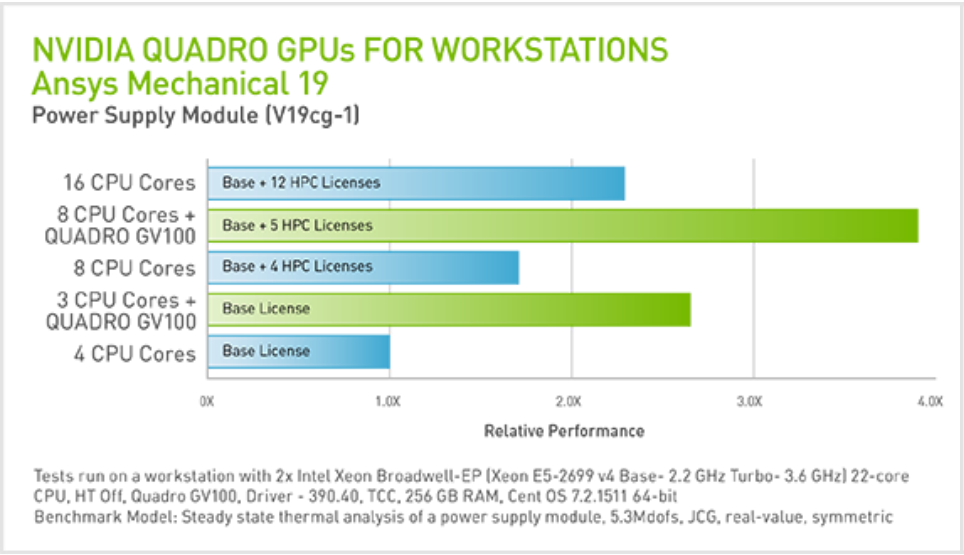

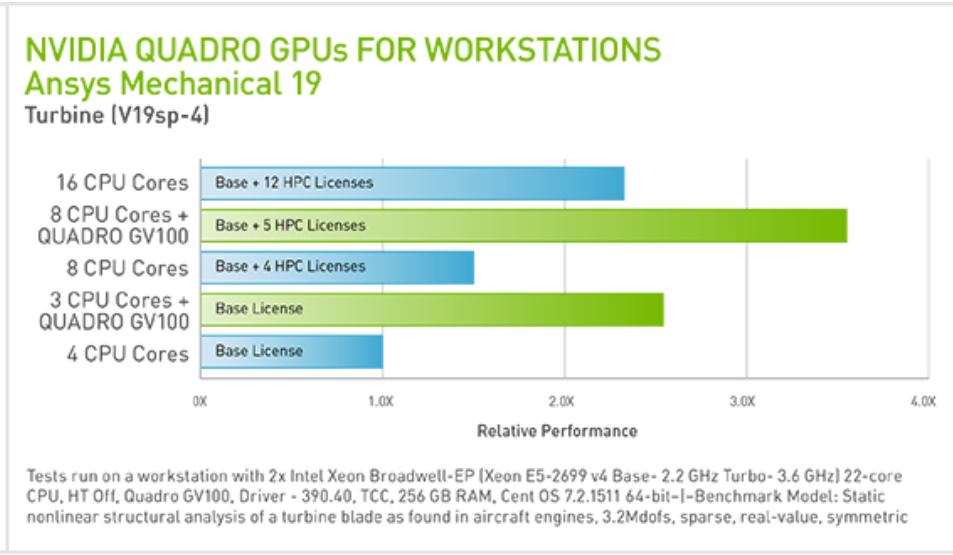

了解互联速度的影响

- **Need fast interconnects to feed fast processors**
  - **Two main characteristics for each interconnect: latency and bandwidth**
  - **Distributed ANSYS is highly bandwidth bound**

```
+--------- D I S T R I B U T E D   A N S Y S   S T A T I S T I C S -----------+

Release: 14.5       Build: UP20120802   Platform: LINUX x64
Date Run: 08/09/2012   Time: 23:07

Processor Model: Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz

Total number of cores available    :   32
Number of physical cores available :   32
Number of cores requested          :    4 (Distributed Memory Parallel)
MPI Type: INTELMPI

  Core                Machine Name   Working Directory
  ------------------------------------------------------
     0                 hpclnxsmc00   /data1/ansyswork
     1                 hpclnxsmc00   /data1/ansyswork
     2                 hpclnxsmc01   /data1/ansyswork
     3                 hpclnxsmc01   /data1/ansyswork

Latency time from master to core    1 =    1.171 microseconds
Latency time from master to core    2 =    2.251 microseconds
Latency time from master to core    3 =    2.225 microseconds

Communication speed from master to core    1 =  7934.49 MB/sec   <- Same machine
Communication speed from master to core    2 =  3011.09 MB/sec   <- QDR Infiniband
Communication speed from master to core    3 =  3235.00 MB/sec   <- QDR Infiniband
```

了解互联速度的影响

- ANSYS Mechanical

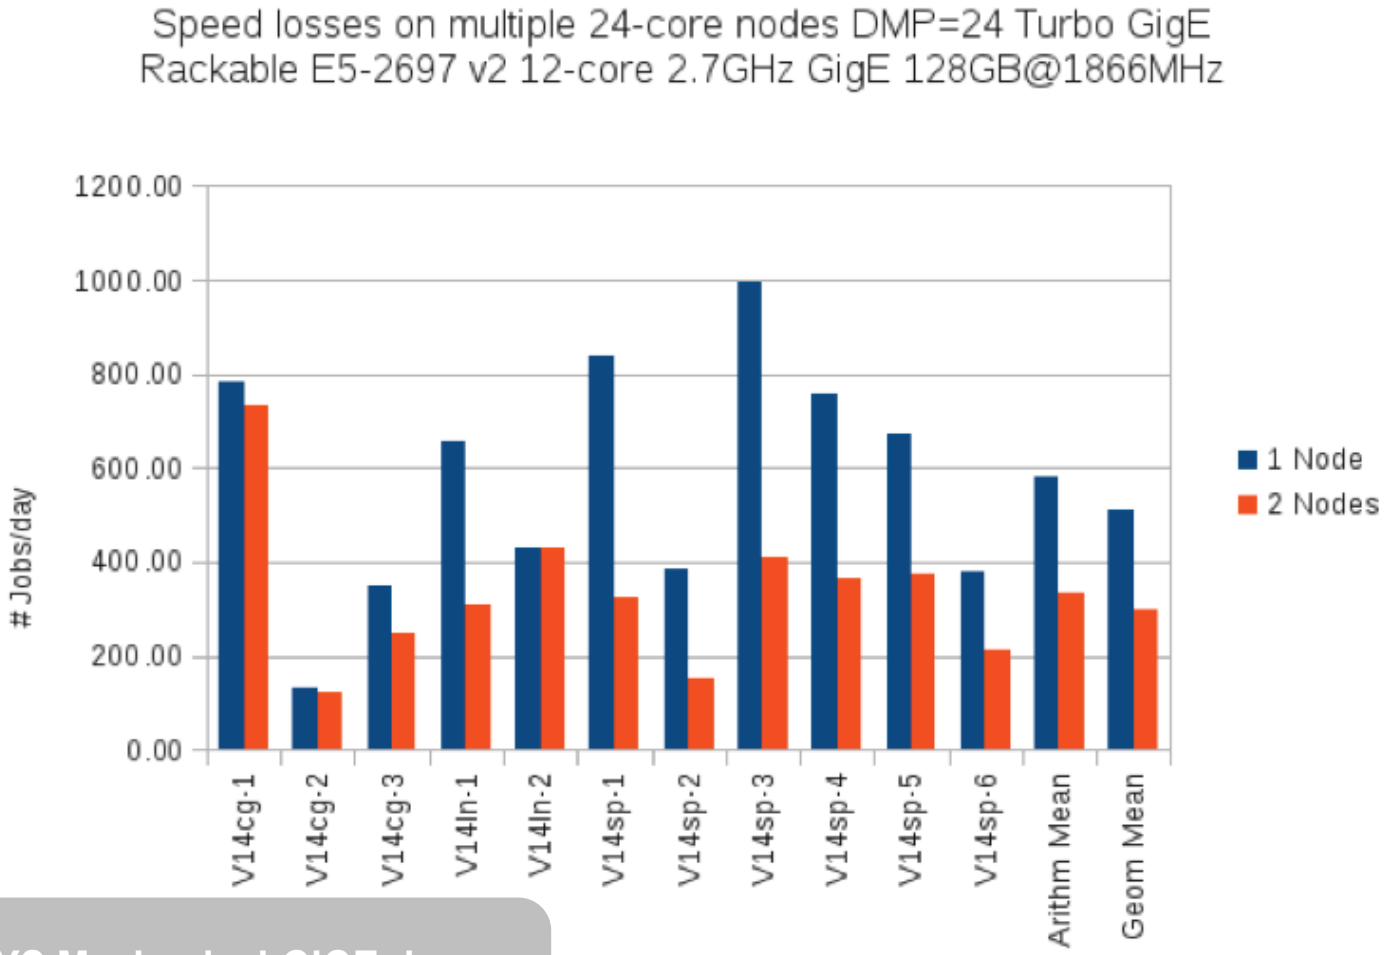

For ANSYS Mechanical GiGE does not scale to more than 1 node!

了解互联速度的影响

- ANSYS Mechanical

V13sp-5 Model

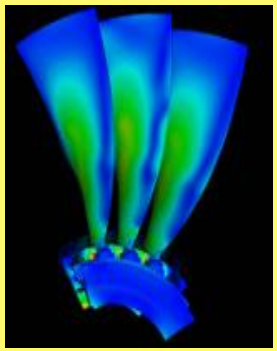

- Turbine geometry
- 2,100 K DOF
- SOLID187 FEs
- Static, nonlinear
- One iteration
- Direct sparse
- Linux cluster (8 cores per node)

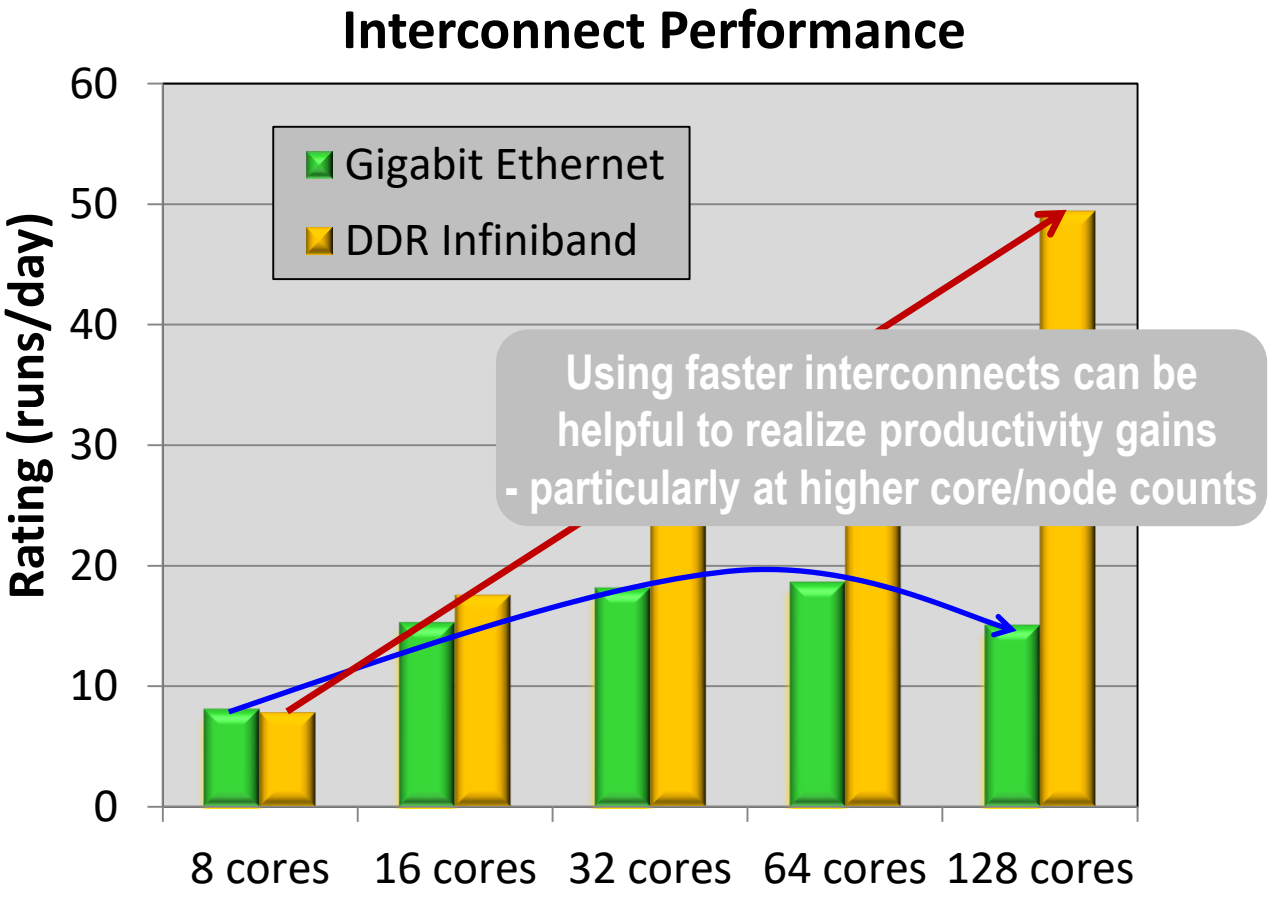

了解存储速度的影响

- ANSYS Mechanical

- **Need fast hard drives to feed fast processors**
  - Check the bandwidth specs
    - ANSYS Mechanical can be highly I/O bandwidth bound
      - Sparse solver in the out-of-core memory mode does lots of I/O
    - Distributed ANSYS can be highly I/O latency bound
      - Seek time to read/write each set of files causes overhead
  - Consider SSDs
    - High bandwidth and extremely low seek times
  - Consider RAID configurations

    RAID 0 – for speed

    RAID 1,5 – for redundancy

    RAID 10 – for speed and redundancy

了解存储速度的影响

- ANSYS Mechanical 18.1

When working directory is assigned to Z Turbo Drive G2 and BMT models for CG solver are used with more than 16 cores, job speeds up by 1.4 times.

When working directory is assigned to Z Turbo DriveG2 and BMT models for SPARSE are used with more than 16 cores, job speeds up by 1.8-2.6 times.

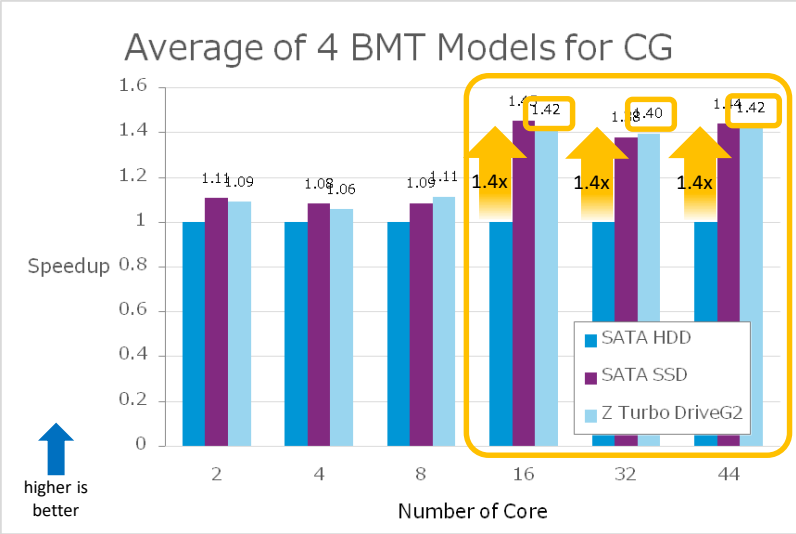

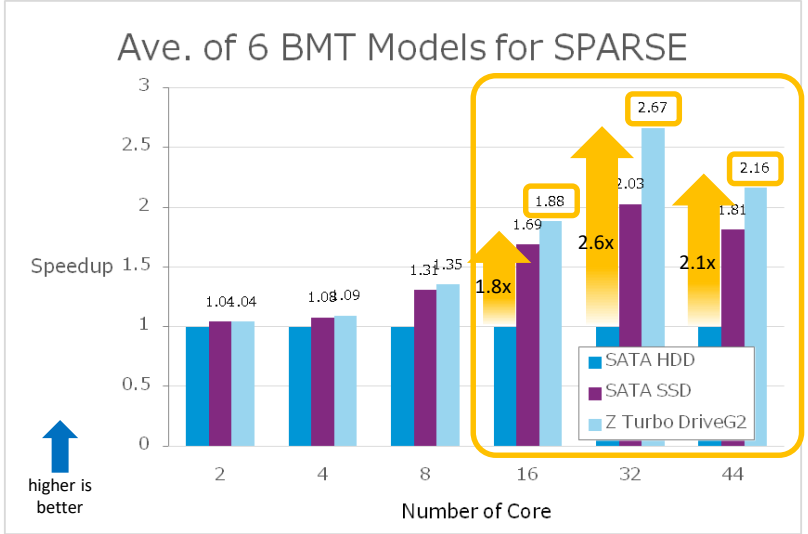

Hardware Configuration:

- **HP Z840 workstation with dual E5-2699v4 (2.2 GHz), 128GBs 2400MHz memory**
- **Optional Storage: Micron SATA SSD No RAID or HP Z Turbo Drive G2 512GB No RAID**

CYBERNET

ANSYS

了解存储速度的影响

- ANSYS Mechanical

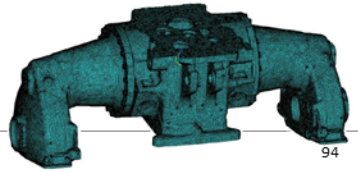

Using faster disks can be helpful to realize productivity gains - particularly at higher core/node counts

小结-硬件配置

- 时钟速度
- 内存带宽
- 互联速度
- GPU加速
- 存储速度：I/O is very important for Mechanical Solver
  - Raid 0 mandatory for multiple disks
  - SSD’s recommended for speed, 15k SAS drives
- Parallel file systems can meet the requirements of both types of solver

典型硬件配置：
ANSYS Mechanical *Starter* Cluster

Hewlett Packard Enterprise

Server Options

– 1 XL1x0r Gen9 head node

– 2-4 ProLiant Xeon nodes, each using 2 processors, in an Apollo 2x00 chassis

  – Up to 28 cores per compute node

  – E5-2690v4 14 core 2.6GHz processors recommended

  – 3-6 480 GB SSD drives (RAID 0)

Options

– 2 NVIDIA K80s (supported on the XL190r only)

Total Memory for the Cluster

– Head and compute nodes: 4 to 8 GBs/ core

Cluster Interconnect

– 10 Gigabit Ethernet or QDR InfiniBand (recommended for jobs using more than 2 nodes)

Operating Environment

– 64-bit Linux, Microsoft (HPC Pack) Server 2012

Workloads

– Suited for Mechanical up to ~80M or ~550M DOFs, depending on the solver used

ANSYS

典型硬件配置：
ANSYS Mechanical *Midsize* Cluster

Hewlett Packard Enterprise

Server Options

- 1 DL360 or DL380 Gen9 head node
- 4-8 ProLiant DL380 Gen9 Xeon server nodes, each using 2 processors (16 cores).
  - 1 NVIDIA K80
  - 2 to 24 internal 600 GB SAS 15K drives or 800 GB SAS SSDs striped RAID 0 per compute node
  - 6x2 TB SAS RAID 0 disk array on the head node
- [OR] 4-8 XL250a Gen9 Xeon server nodes (Apollo 6000), each using 2 processors (16 cores)
  - Up to 2 NVIDIA K80s
  - 6 internal SAS 15K drives or 800 GB SAS SSDs per compute node (suitable for nonlinear jobs > = 2M DOF
- E5-2667v4 3.2 GHz 8 core processors recommended

Total Memory for the Cluster

- Head node: 16 GB/ core
- Compute nodes: 6-12 GB/ core on each remaining node

Cluster Interconnect

- EDR InfiniBand for larger clusters

Operating Environment

- 64-bit Linux or Microsoft HPC Server 2008

Workloads

- 512-2048 GB RAM configurations will handle up to ~8 simultaneous running ANSYS “mega models” of ~55 – 230M or ~580 DOFs, depending on the solver used

ANSYS

应用案例

ANSYS Mechanical支持GPU加速

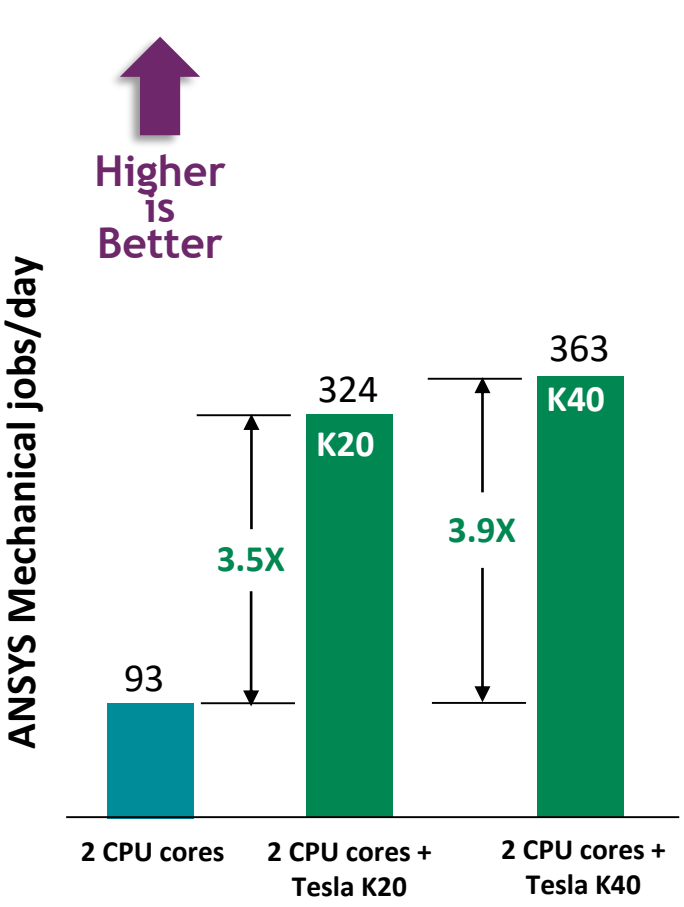

Simulation productivity (with a HPC license)

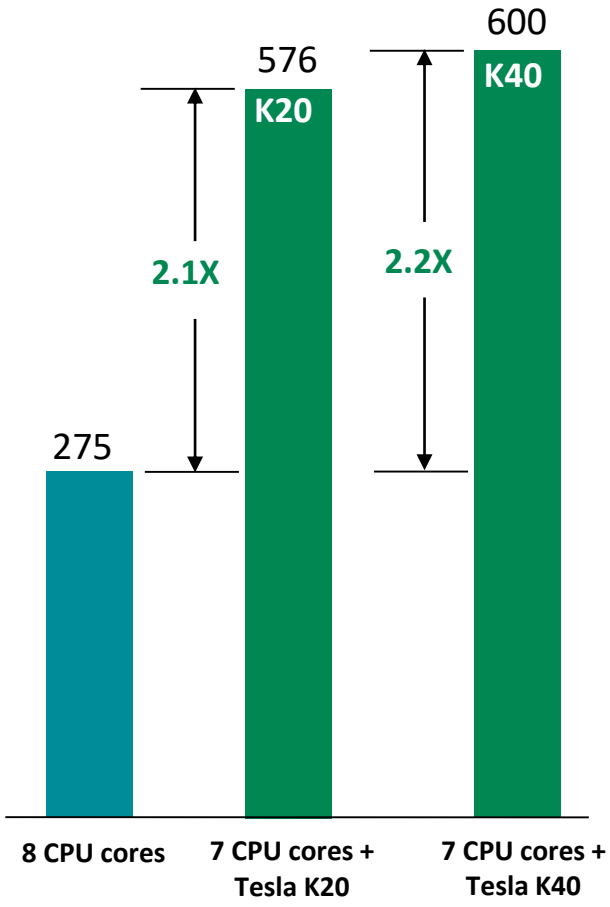

Simulation productivity (with a HPC Pack)

V14sp-5 Model

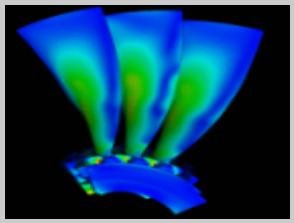

- Turbine geometry
- 2.1 million DOF
- SOLID187 elements
- Static, nonlinear analysis
- One iteration
- Sparse direct solver

Distributed ANSYS Mechanical 15.0 with Intel Xeon E5-2697 v2 2.7 GHz CPU; Tesla K20 GPU and a Tesla K40 GPU with boost clocks.

应用案例

- Total time to complete one run

Model: NNG Container
Nonlinear Static analysis
    Solid elements, nonlinear geometry and material
    1 148 626 nodes (3.4 MDOF)

Benchmark completed on:
Intel(R) Xeon(R) CPU E5-2660v3 @ 2.60GHz
20 core cluster nodes
128 GB memory available/node
SUSE Linux Enterprise Server 11.3

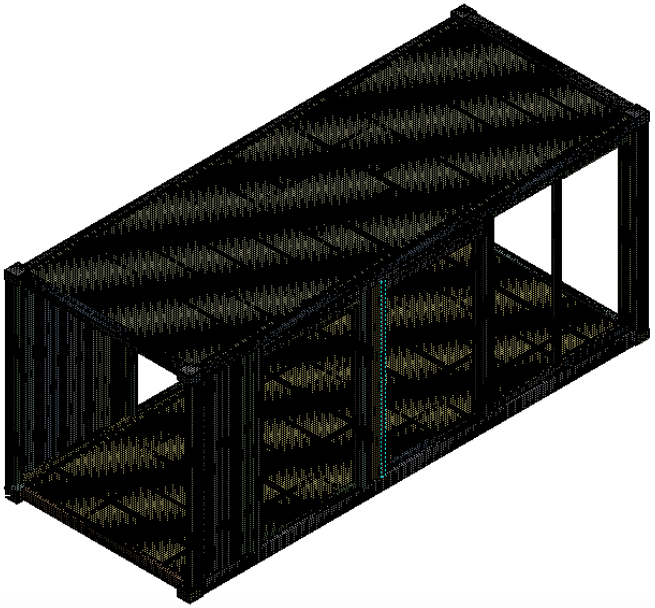

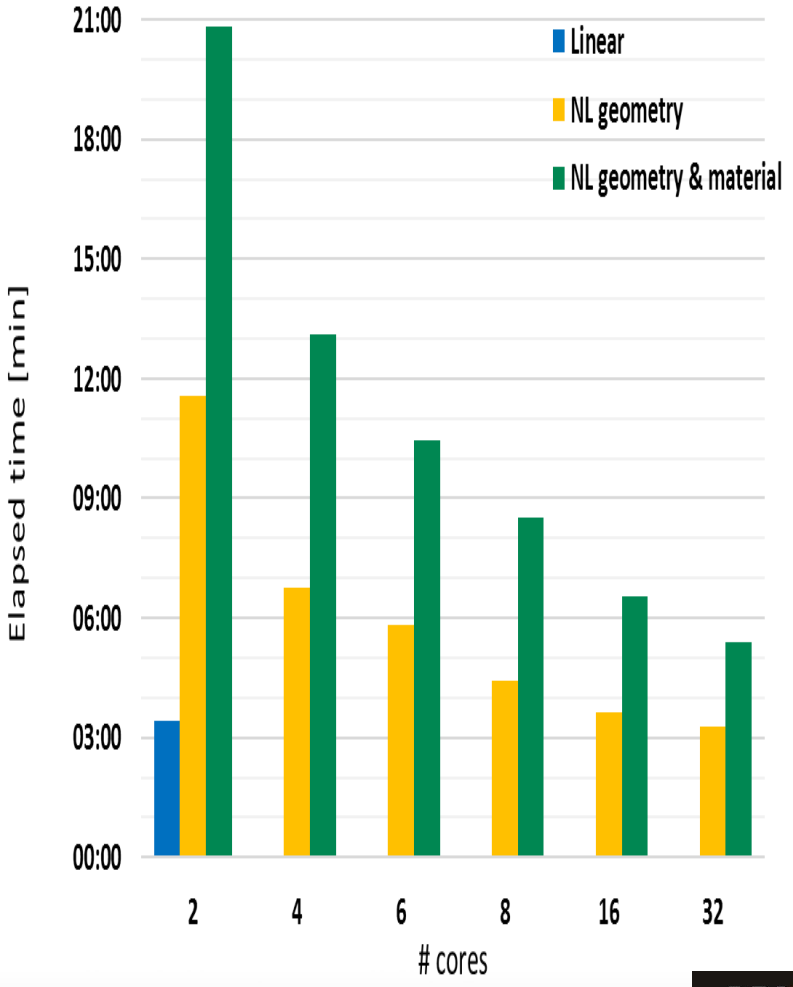

感谢聆听！

ANSYS-China

ANSYS中国