Accelerating Simulation & Analysis with Hybrid GPU Parallelization and Cloud Computing

Devin Jensen

August 2012

Altair Knows HPC

Altair

Altair is the only company that:

makes HPC tools…

PBS Works

develops HPC applications…

HyperWorks

…and **uses** these to solve real HPC problems

Altair ProductDesign

500 Altair engineers worldwide

use HPC every day *for*

real-world modeling

& simulation

Innovation Intelligence®

Altair

26+

Years of Innovation

40+

Offices in 18 Countries

1500+

Employees Worldwide

Customers

Altair

Automotive	Aerospace	Heavy Equipment	Government	Life/Earth Sciences	Consumer Goods	Energy
Chrysler, Magna, Mazda, GM, FAW, Toyota, Peugeot, Tata, SAIC, Sumitomo, Fiat, Hitachi, Volvo, Aisin, Bridgestone, Hyundai, Honda, BMW, DFM, Renault, Mitsubishi, Nissan, VW, Denso	BAE Systems, Boeing, GE Aircraft Engines, General Dynamics, EADS, Bombardier, Honeywell, Lockheed Martin, Lufthansa, Finmeccanica, Raytheon, Rolls-Royce, United Technologies, Northrop Grumman	Alstom, John Deere, Caterpillar, JFE, Komatsu, Tata Steel, Hitachi, Bombardier, Nippon Steel, CRCC 中国铁建, Norinco, Kobelco, Mitsubishi Heavy Industries, Oshkosh, Itochu, Doosan, CSIC, ArcelorMittal, Baosteel 宝钢, Hyundai Heavy Industries, SNCF, VW	Argonne National Laboratory, ARL, AIST, U.S. Air Force, DLR, Oak Ridge National Laboratory, NASA, National Research Council Canada, Sandia National Laboratories, NAVSEA, Los Alamos National Laboratory	BASF The Chemical Company, Adecco, Abbott, Bayer, Dow, Fresenius Medical Care, Merck, Vale, Alcoa, sanofi aventis, lyondellbasell, accenture High performance. Delivered., SMFG, Sabic	acer, Compal, Huawei, Canon, Henkel, HP, 3M, Ricoh, Samsung, Intel, IBM, Fujitsu, RIM, Tesco, Xerox, LG, Apple, Quanta Computer, Lenovo, Bosch, NEC, Philips, Microsoft, DuPont, Suzuki, Emerson, Motorola, Kirin, Toshiba, Panasonic, Yamada, Flextronics, Sony Ericsson, P&G, Wistron, Nokia Connecting People, Johnson Controls, Walt Disney Pictures, Unilever, Fujifilm	ABB, Chevron, Shell, Schneider Electric, Petrobras, ConocoPhillips, Total, SK holdings, GS, Schlumberger, Reliance Industries Limited, Valero, ExxonMobil, Enel, Siemens, Evonik, Idemitsu, Petronas, ThyssenKrupp

3,200+ customers worldwide

Hardware Trends and Software Evolution

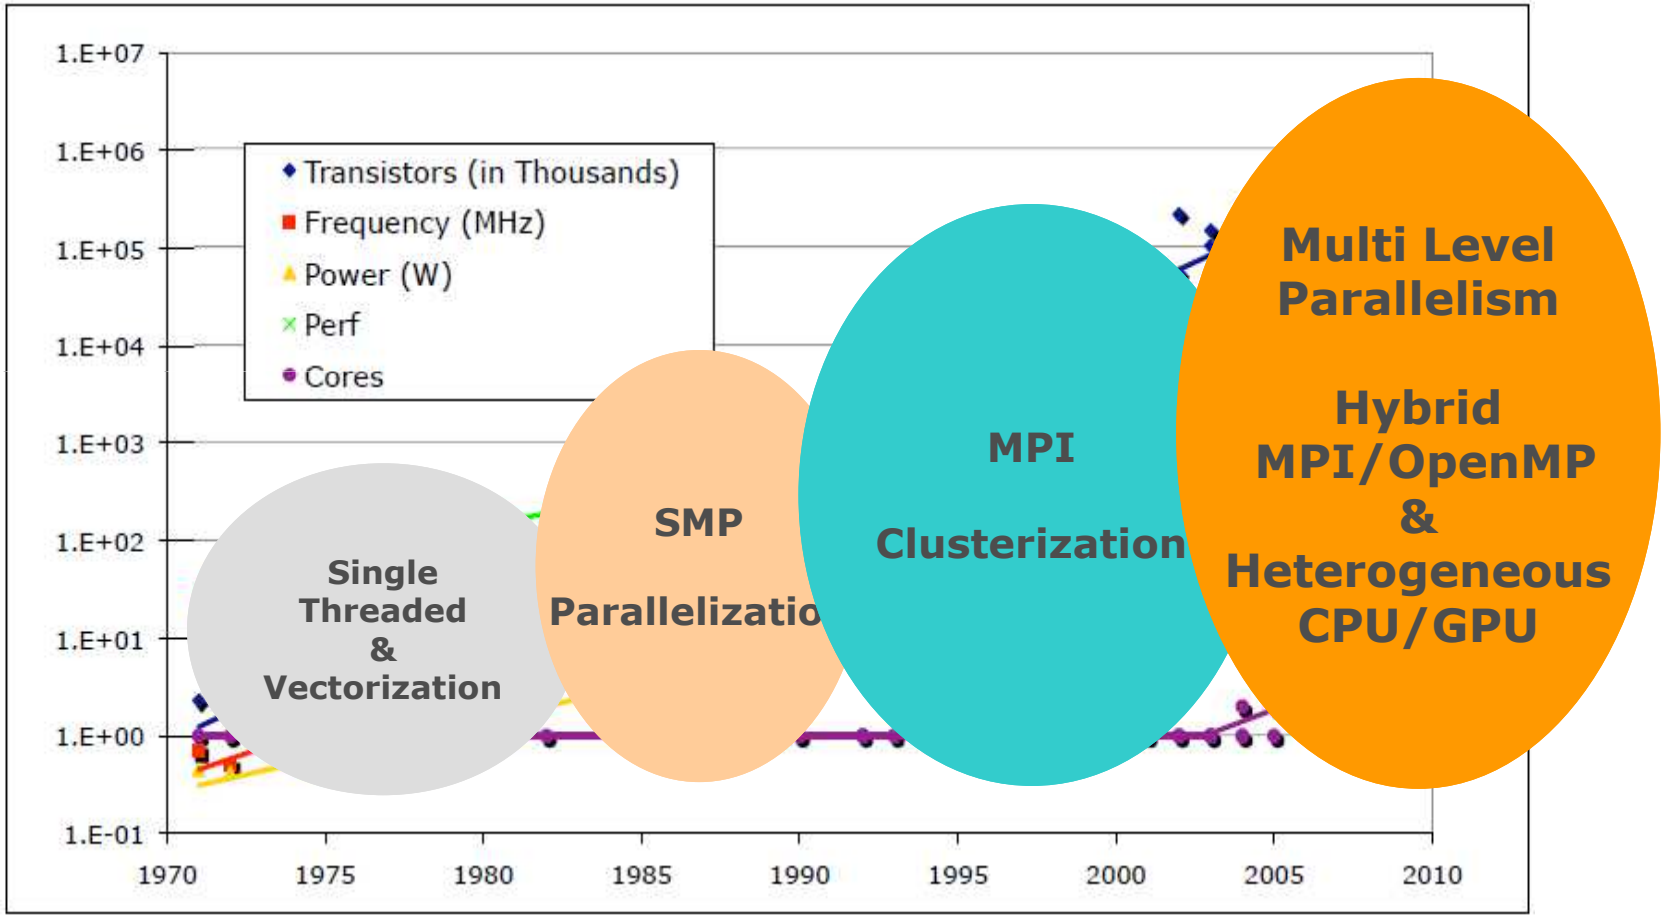

Source : Intel Corp.

Motivations to use GPU for Solvers

HyperWorks

- Cost effective solution
- Power efficient solution

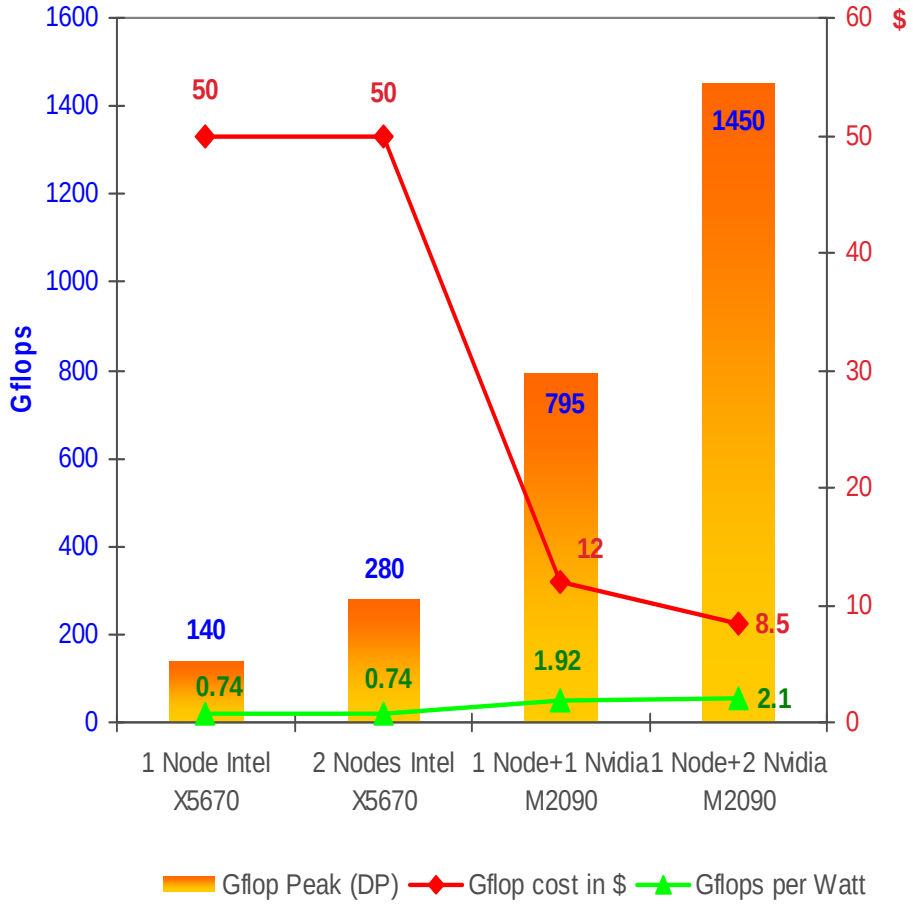

How much of the peak can we get?

- Which part of the code is best suited?
- How much coding effort is required?
- What will be the speedup for my application?

NVIDIA

- **Ecosystem based on CUDA (+ OpenACC with CRAY, PGI and CAPS)**
  - Any NVIDIA graphic card supports CUDA
  - Strong market presence
- **Products**
  - Tesla and Quadro GPUs based on the Fermi Architecture
  - Announcement of Kepler Architecture, ~2X faster than Fermi, HyperQ, virtualization,…

AcuSolve GPU Porting on NVIDIA CUDA

HyperWorks

- High performance computational fluid dynamics software (CFD)
- The leading finite element based CFD technology
- High accuracy and scalability on massively parallel architectures

S-duct
80K Degrees of Freedom
Hybrid MPI/OpenMP for Multi-GPU test

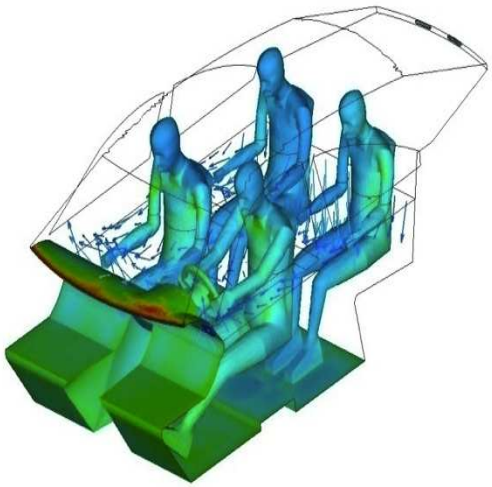

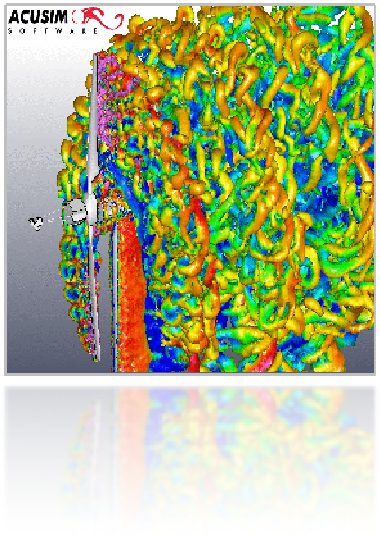

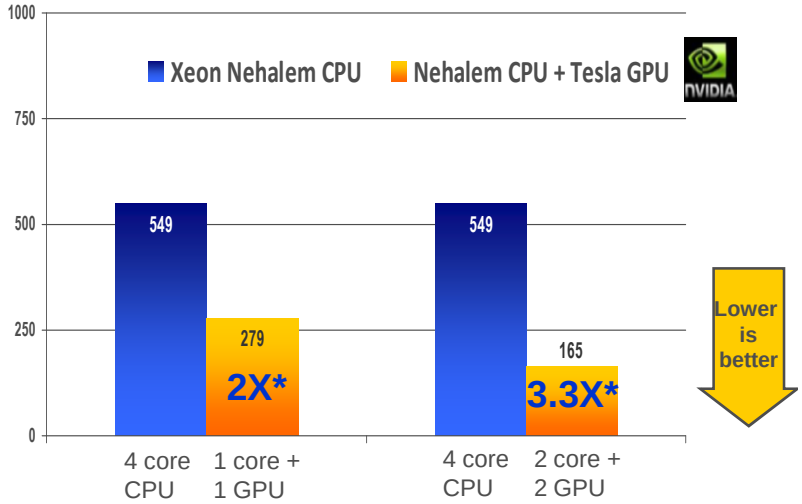

*Performance gain versus 4 core CPU

RADIOSS Porting on GPU

HyperWorks

- **Assess the potential of GPU for RADIOSS**
- **Focus on Implicit**
  - Direct Solver
    - ✓ Highly optimized compute intensive solver
    - ✓ Limited scalability on multicores and cluster
  - Iterative Solver
    - ✓ Ability to solve huge problems with low memory requirement
    - ✓ Efficient parallelization
    - ✓ High cost on CPU due to large number of iterations for convergence
- **Double Precision required**
- **Integrate GPU parallelization into our Hybrid parallelization technology**

RADIOSS Porting on GPU & Accelerators

HyperWorks

- **CUDA official version planned for HW12**
  - RADIOSS implicit direct solver and iterative solver available under Linux
  - Support for NVIDIA Fermi and Kepler (K20 Q1’13)
  - Altair RADIOSS presentation at GTC2012

- **OpenCL Beta version**
  - Only RADIOSS Implicit Iterative Solver available under Linux and Windows
  - Support for AMD FirePro
  - Altair RADIOSS presentation at AFDS2011

RADIOSS Direct Solver GPU Porting

HyperWorks

- **CUBLAS (DGEMM) – perfect candidate to speed up update module**
  - Frontal matrix could be too huge to fit in GPU memory
  - Frontal matrix could be too small and thus inefficient w/ GPU
  - Data transfer is not trivial
  - Only the lower triangular matrix is interesting

$$A = \begin{pmatrix} B & V^t \\ V & C \end{pmatrix}$$

```
do j = 1, N
    assemble(A(j))
    factor(j)
    update(j)
end
```

Non-pivoting: Cholesky
Pivoting: LDLT

- **Pivoting is required in real applications**
  - Pivot searching is a sequential operation
  - Factor module has limited parallel potential
  - It could be as expensive as “update” module in extreme case

$$= \begin{pmatrix} L_B & 0 \\ VL_B^{-t} & I \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & C - VB^{-1}V^t \end{pmatrix} \begin{pmatrix} L_B^t & L_B^{-1}V^t \\ 0 & I \end{pmatrix}$$

- **Asynchronous computing**
  - Overlap the computation
  - Overlap the communication

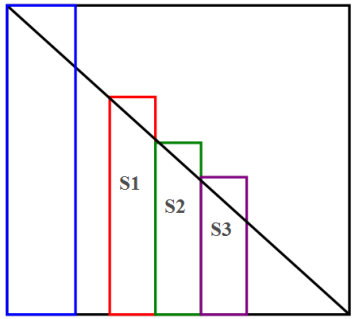

$$U = LDL^T$$

RADIOSS Direct Solver – Non Pivoting Case

HyperWorks

Linear static – Non Pivoting Case	
Benchmark	2.8 Millions of Degrees of Freedom
Platform	Intel Xeon X5550, 4 Core, 48GB RAM NVIDIA C2070, CUDA 4.0 MKL 10.3, RHEL 5.1

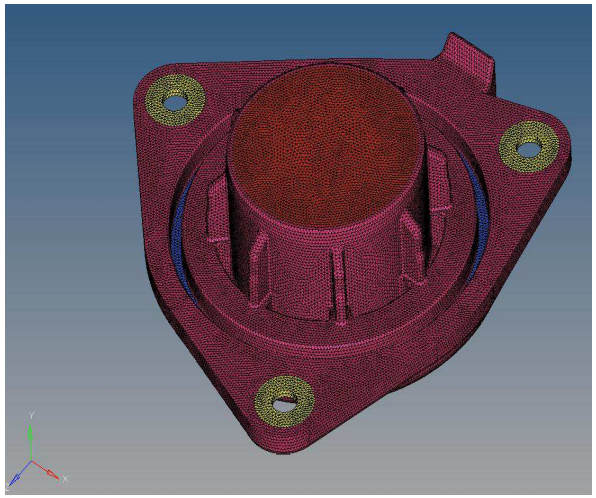

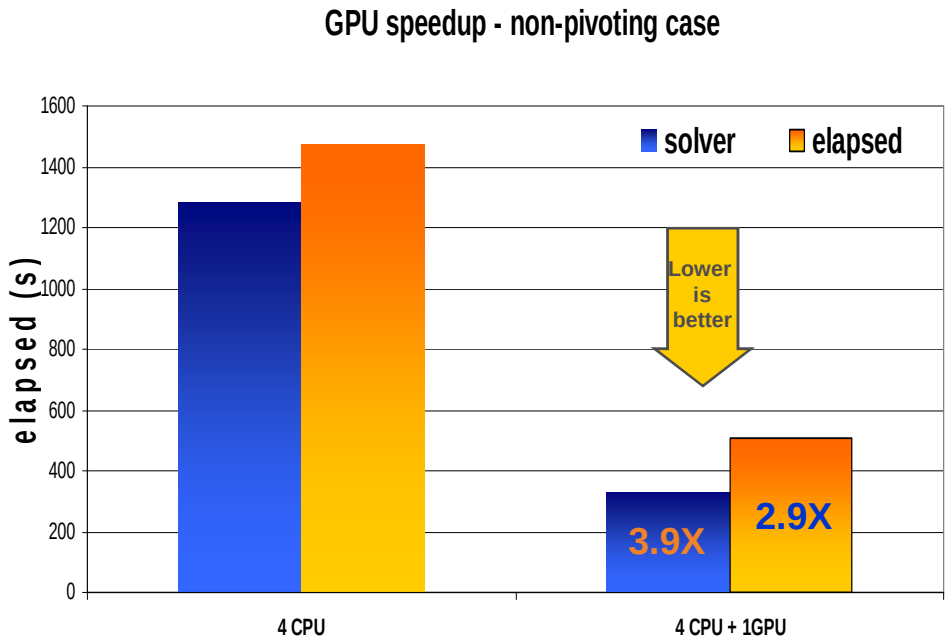

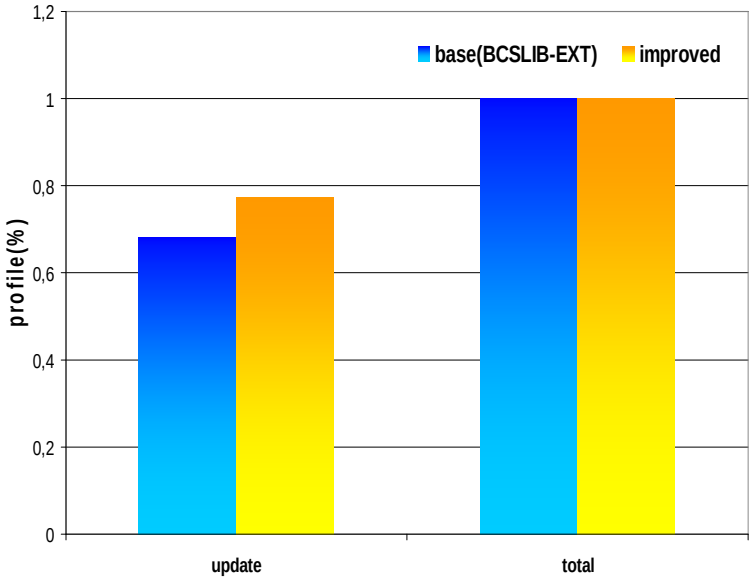

RADIOSS Iterative Solver (Implicit)

HyperWorks

- Preconditioned Conjugate Gradient (PCG) solves iteratively $M^{-1} . (Ax – b) = 0$
  - Convergence speed depends on preconditioner M
  - Low memory consumption
  - Efficient parallelization: SMP, MPI

- Porting under CUDA
  - Few kernels to write in Cuda ⇨ Sparse Matrix Vector
  - Use of CUBLAS ⇨ DAXPY & DDOT

- Extend Hybrid to multi GPUs programming
  - MPI to manage communication between GPU
  - Portions of code not GPU enabled benefit from OpenMP parallelization
  - Programming model expandable to multi nodes with multiple GPUs

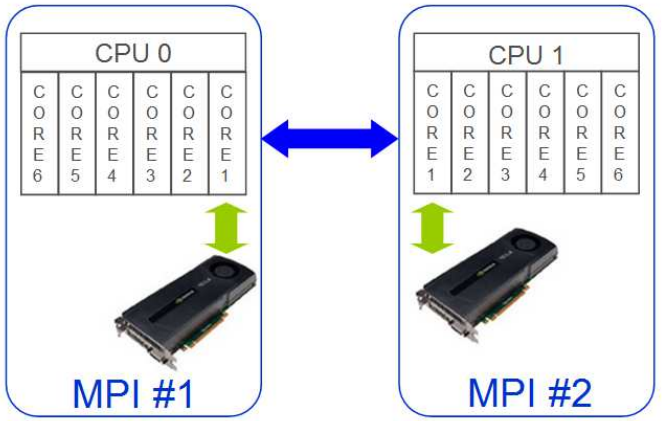

RADIOSS PCG – Linear Benchmark #2

HyperWorks

Linear Problem #2 **Hood of a car with pressure loads** **Refined model** **Compute displacements and stresses**	
Benchmark	**2.2 Millions of Degrees of Freedom** **62 Millions of non zero** **380000 Shells + 13000 Solids + 1100 RBE3** **5300 iterations**
Platform NVIDIA	**NVIDIA PSG Cluster – 2 nodes with:** **Dual NVIDIA M2090 GPUs** **CUDA v3.2** **Intel Westmere 2x6 X5670@2,93Ghz** **Linux RHEL 5.4 with Intel MPI 4.0**
Performance	**Elapsed time decreased by up to 13X**

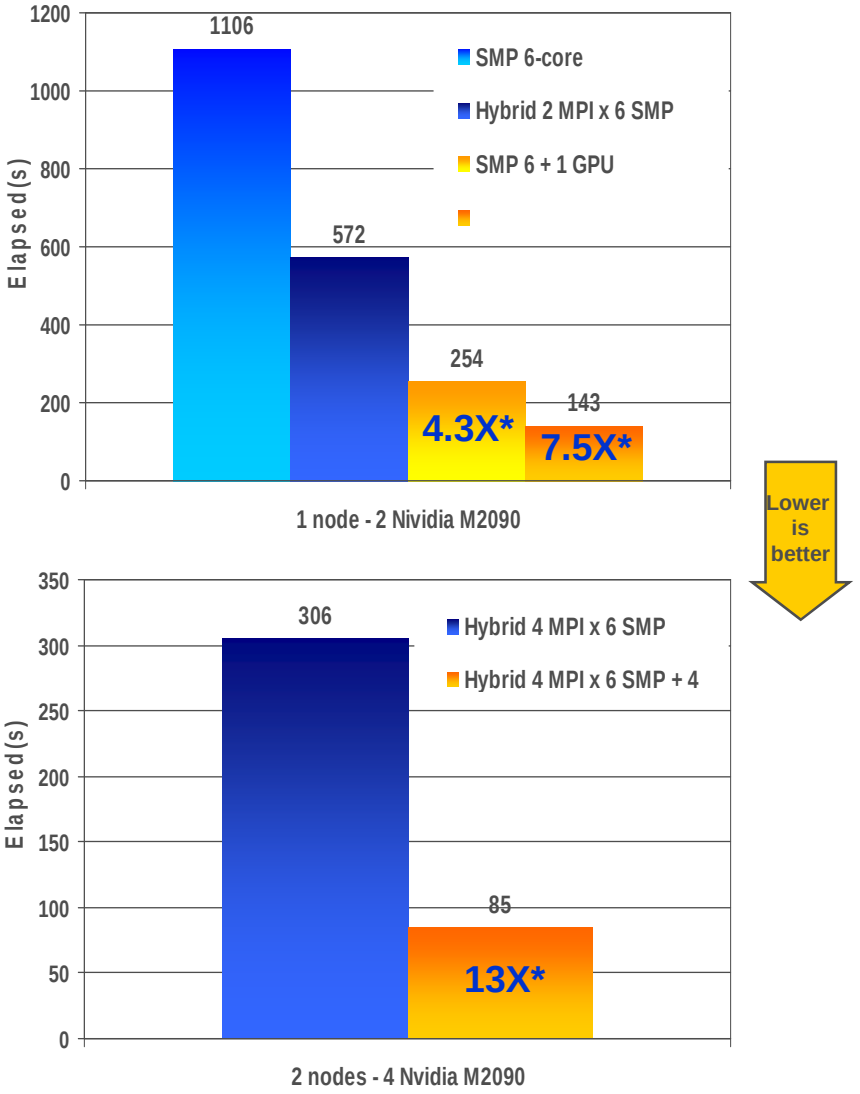

Multi GPUs performance is better for bigger models

**Performance gain versus SMP 6-core*

GPU For Explicit ?

HyperWorks

- **Explicit is very difficult to port on GPU**
  - Data movement between CPU and GPU current bottleneck
  - Require full memory allocated on GPU and therefore full code ported
    - ✓ Huge development cost
    - ✓ Challenge of multiple code bases
  - Potential gain remains small with current technology approach
    - ✓ Data locality and cache reuse is not possible on GPU
    - ✓ Memory size limitation
- **Follow hardware evolution**
  - AMD APU (CPU+GPU)
  - Intel Xeon Phi
  - NVIDIA through OpenACC initiative

Altair Solvers on NVIDIA GPUs

HyperWorks

- RADIOSS implicit direct & iterative solvers have been successfully ported on NVIDIA GPU using CUDA
- Adding GPU improves performance of Altair solvers significantly
- For iterative solver, Hybrid MPP allows to run on multi GPU card workstation and GPU cluster with good scalability and enhanced performance
- GPU support for implicit solvers is planned for HyperWorks 12

HyperWorks 3D display

- OpenGL stereoscopic display in HyperMesh & HyperView
- Supported by NVIDIA graphics cards
- Compatible with NVIDIA 3D Vision glasses
- Available in HyperWorks 12.0

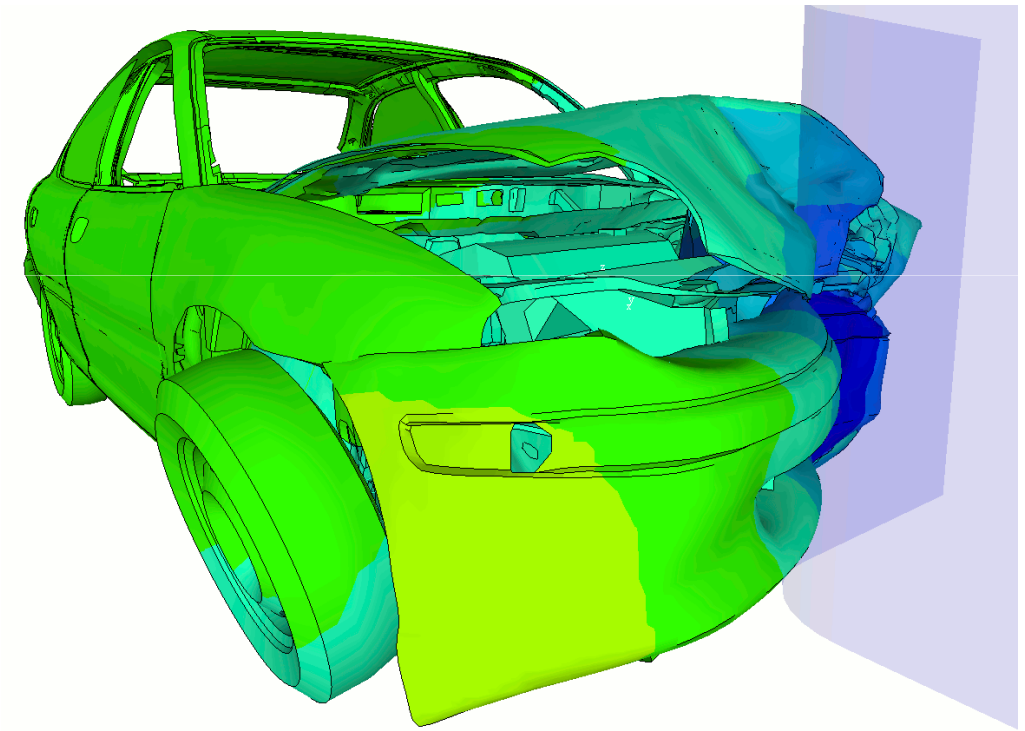

HyperView Performance Gains From Shaders

HyperWorks

Gains in graphics performance – 2x to 6x improvement in animation frame rate, model rotation, & pan

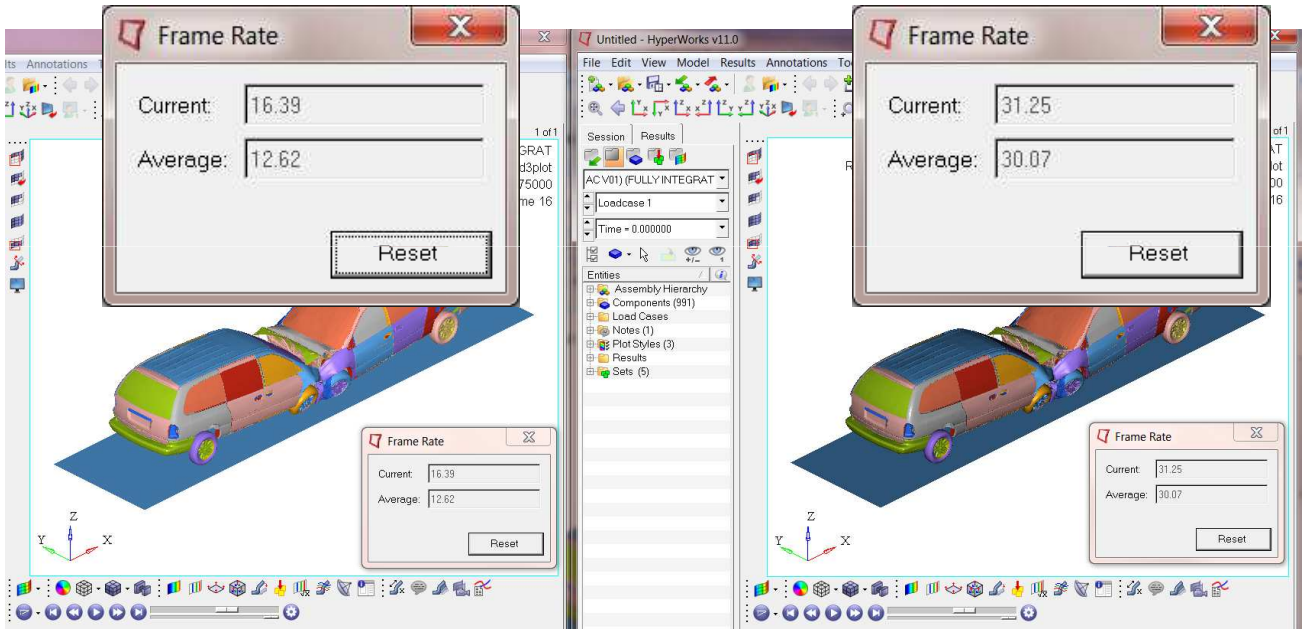

Licensing Model for Solvers

- Licensing has been set to encourage GPU testing
- One GPU is considered as one additional core only

Where does PBS Professional manage NVIDIA GPUs?

PBS Works

Tokyo Institute of Technology – TSUBAME 2.0

- Reliability
- Robust advance reservations
- GPU support
- Mixed architecture support (Linux, Windows)

~100,000 simultaneous jobs

~2,000 users

2.4 Petaflops
17,984 cores
4,200 GPUs
80 TB mem
174 TB SSD
7 PB disk

PBS Professional offers THE best scalability in the industry…

– NEC

TOP500 SUPERCOMPUTER SITES

HP's biggest Top500 machine

*Decreed as one of the "**Greenest Production Supercomputer in the World**" on the June 2012 Green 500 List, for achieving a high level of energy-efficiency*

Purpose-Built for High Performance Computing

Altair

IDC study (published 2009) on job scheduling ranked PBS Professional #1

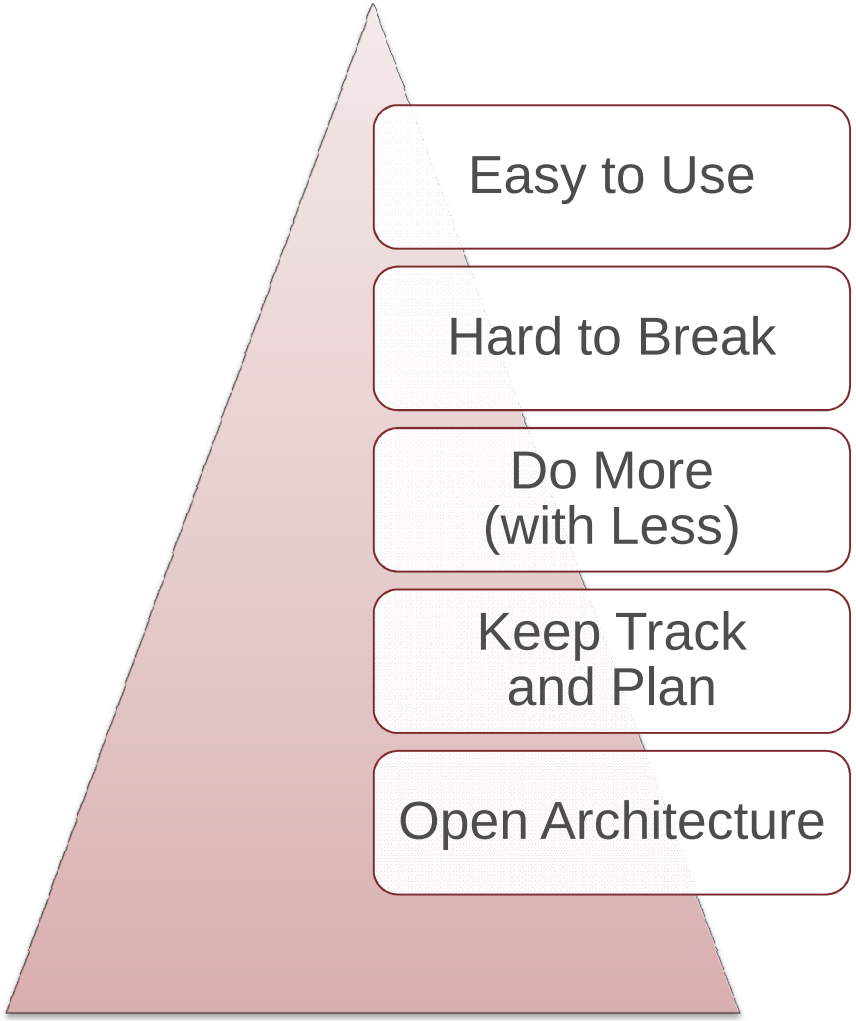

5 Strategic Pillars

Compute Manager

PBS Professional®

PBS Analytics™

Altair

Structure our Thoughts

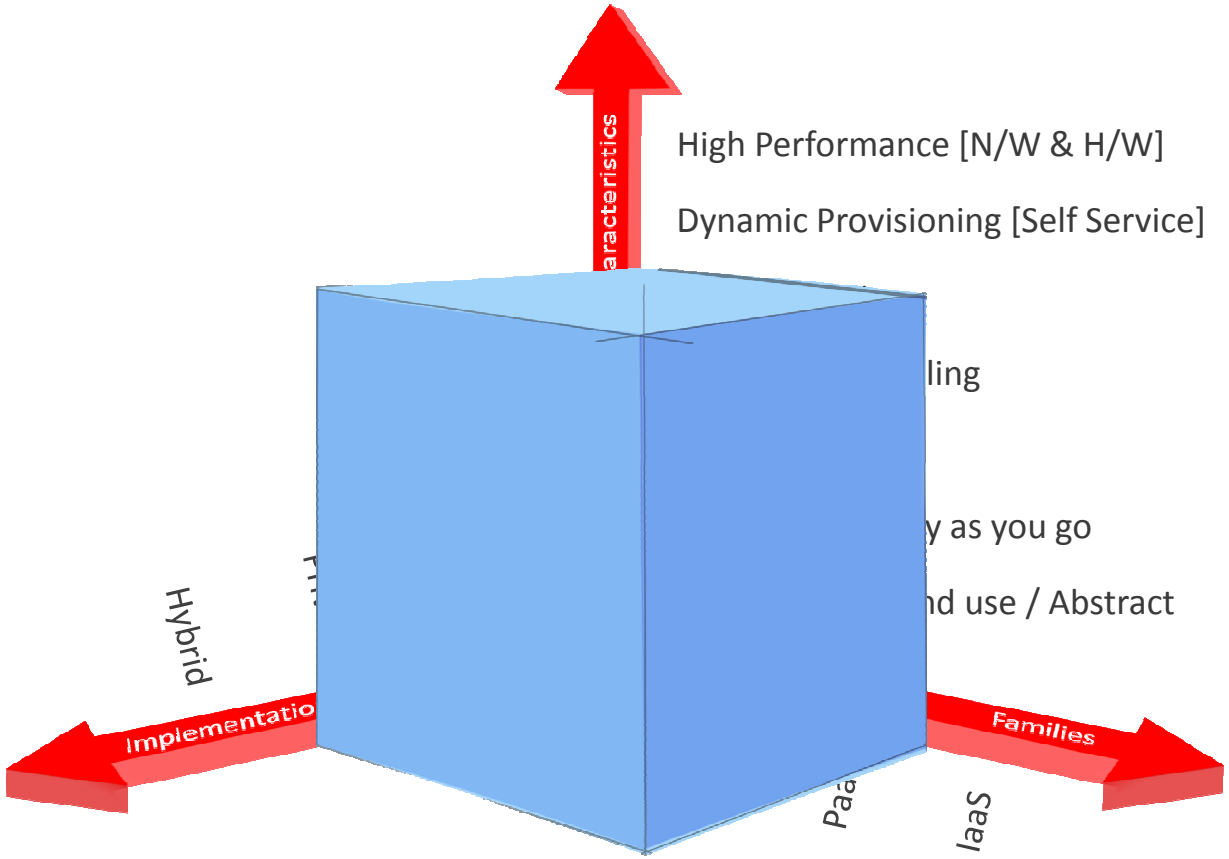

End Users Needs

Focus on My Task / Project → ***Deliver Results!***

Does have

Workstations

“Limited” Access to Software

Ideas, Innovation, Agility

Does NOT have

Expertise in IT Infrastructure

Sufficient compute resources

Financial / Abundance of Money

End Users Needs

Ease of Access: All you need is just a Browser
to access HPC & Domain specific resources

Ease of Use: All you need is a browser based application
hiding the complexity

Handle Big Data: All you need is a way to visualize the managed
big data again through a browser based application

User Interface → **Portals**

Altair Enterprise Portals

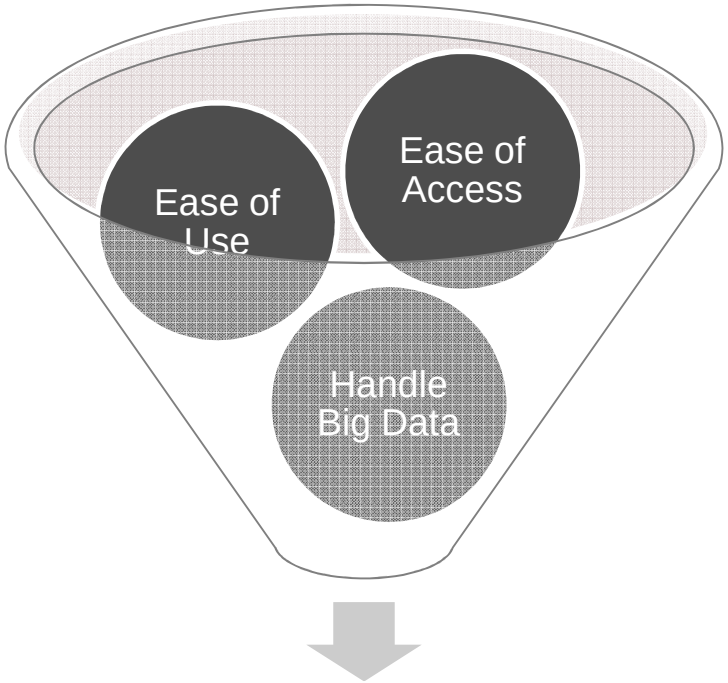

HyperWorks Enterprise Applications

- Compute Manager
- Display Manager
- Simulation Manager
- Process Manager
- Performance Manager

COMPUTE MANAGER

Ease of Use | Ease of Access | Handle Big Data

Compute Manager

Let’s See

Ease of Use | **Ease of Access** | Handle Big Data

Compute Manager

Let's See

Ease of Use | Ease of Access | **Handle Big Data**

Altair

HyperWorks Enterprise

Complete Simulation Lifecycle Management

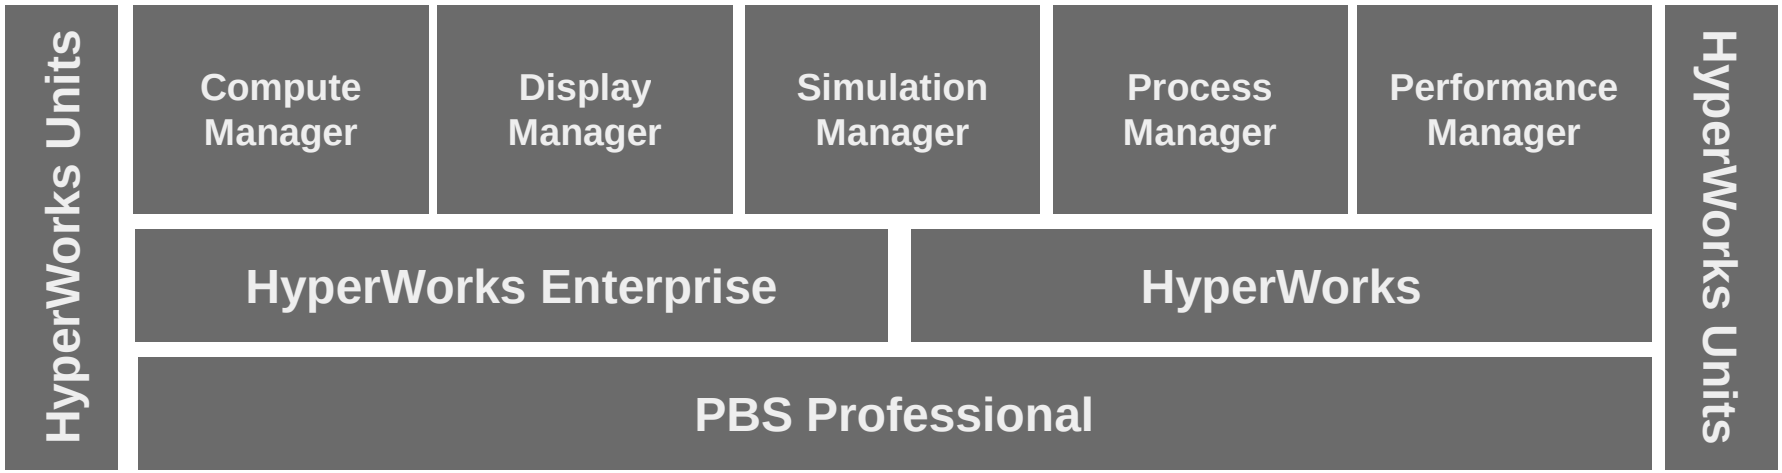

It is "HyperWorks in the Cloud"

Implementation Models

HYPERWORKS ON-DEMAND

A Public Cloud run by Altair

HyperWorks On-Demand: The Basics

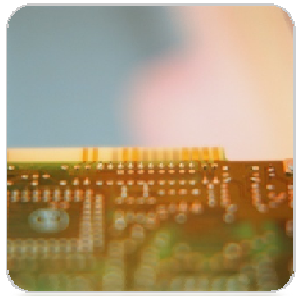

State-of-the-art HPC computation nodes

Scalable up to 10,000 cores

QDR Non-blocking InfiniBand

Physical and cyber security

Resource Provisioning

Scheduling

Security Framework

Licensing Framework

Remote Visualization Framework

Notification Framework

Portal & Mash up Framework

Collaboration Framework

Powered by PBS Professional & HyperWorks Enterprise

HyperWorks Solvers

HyperWorks Desktop[1]

Altair Partner Alliance[1]

1: coming soon

IaaS + PaaS + SaaS = HWaaS

Altair

The High-Powered Altair Data Center

- Scalable, modular facility located in Troy, MI
- Can be easily extended to support up to more than 10,000 cores
- Incorporates extensive physical and cyber security measures
- Extremely high compute-power density
- A state-of-the-art facility for HPC simulation computing

Accelerating your Gateway to HPC Cloud Computing

Altair

HyperWorks On-Demand

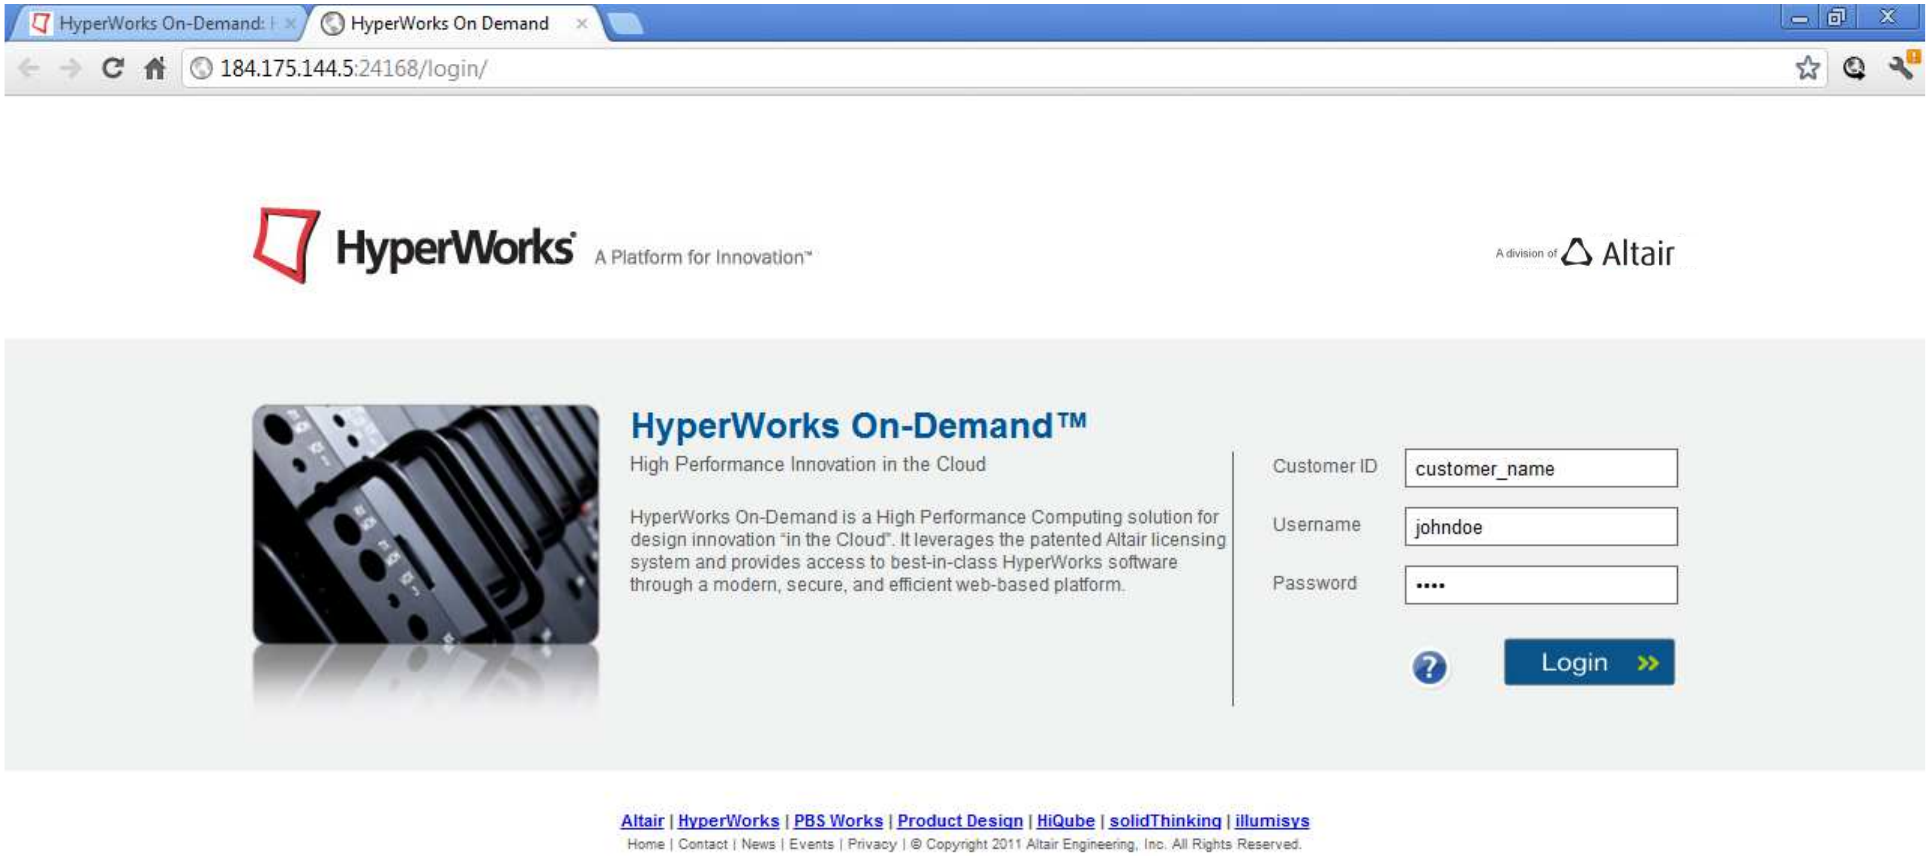

Altair's user portals is the gateway to HyperWorks On-Demand

Altair Knows HPC

Altair

Altair is the only company that:

makes HPC tools…

PBS Works

develops HPC applications…

HyperWorks

…and **uses** these to solve real HPC problems

Altair ProductDesign

500 Altair engineers worldwide

use HPC every day *for*

real-world modeling

& simulation

For More Information

Altair

- **Contact Us**
  - Devin Jensen
  - Director, Americas PBS Field Operations
  - jensen@altair.com
- **Visit Us Online**
  - www.altair.com
  - www.altairhyperworks.com
  - www.pbsworks.com

Thanks, NVIDIA

Thank YOU!