SGI®

Many core computing Solutions

Gábor Lehoczki
Silicon Computers Kft.
leho@silicon.hu

Nvidia Tesla series

Features	M2075	M2090	K10	K20	K20X
Number and Type of GPU	1 Fermi GPU	1 Fermi GPU	2 Kepler GK104s	1 Kepler GK110	
Peak double precision floating point performance	515 Gigaflops	665 Gigaflops	190 Gigaflops (95 Gflops per GPU)	1.17 Tflops	1.31 Tflops
Peak single precision floating point performance	1030 Gigaflops	1331 Gigaflops	4577 Gigaflops (2288 Gflops per GPU)	3.52 Tflops	3.95 Tflops
Memory bandwidth (ECC off)	150 GBytes/sec	177 GBytes/sec	320 GB/sec (160 GB/sec per GPU)	208 GB/sec	250 GB/sec
Memory size (GDDR5)	6 GB	6 GB	8 GB (4 GB per GPU)	5 GB	6 GB
CUDA cores	448	512	3072 (1536 per GPU)	2496	2688

Intel® Xeon Phi™ Coprocessor

	5110P	3120A	3120P	7120P	7120X	5120D
Code Name	Knights Corner					
# of Cores	60	57	57	61	61	60
Clock Speed	1.053 GHz	1.1 GHz	1.1 GHz	1.238 GHz	1.238 GHz	1.053 GHz
Peak double precision floating point performance						1.011 TFlops
Peak single precision floating point performance						2.022 TFlops
Cache	30 MB	28.5 MB	28.5 MB	30.5 MB	30.5 MB	30 MB
Max TDP	225 W	300 W	300 W	300 W	300 W	245 W
Max Memory Size	8 GB (16 channels)	6 GB	6 GB	16 GB	16 GB	8 GB
Max Memory Bandwidth	320 GB/s	5 GB/s	5 GB/s	5.5 GB/s	5.5 GB/s	5.5 GB/s

SGI Development Suite

- For Linux software development
- Supports 2-4 developers

- Technical Computing Performance for SGI Systems
- Consists of SGI Accelerate, SGI MPI, SGI REACT, SGI UPC
- Tech support via SGI Support contract

- Software development tool suite tuned for application performance and code robustness
- Consists of C++ & Fortran compilers, Math Kernel Library, Integrated Performance Primitives, Threading Building Blocks
- Includes 1 yr tech support

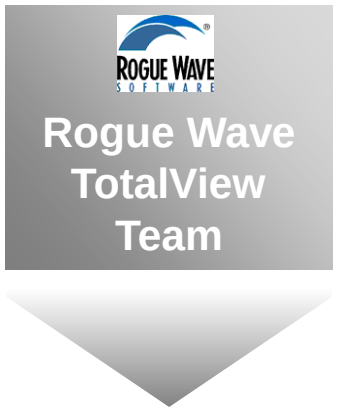

- Dynamic source code, thread, and memory debugging for C, C++ and Fortran HPC applications
- Consists of TotalView, MemoryScape, ReplayEngine, and CUDA debugging
- Includes 1 yr tech support

4

Intel® Xeon Phi™ compared to a GPU

- Intel® Xeon Phi™ is a multi-core architecture with the following characteristics :
  - based on Pentium 4 in-order cores with a vector extension
  - OpenMP programming
  - “UV on a chip”
  - PCIe x16 Gen 2 (~ 6.7 GB/s)
  - Ease-of-Use
- NVIDIA Fermi GPU use streaming multi-processors with currently 32 SIMD engines
  - A very high number of threads (10s of thousands) hides latency to memory
  - Minimal context switching time between threads
  - HW thread dispatcher
  - PCIe x16 Gen (~ 11.3 GB/s) [Kepler 10]

GPU-ACCELERATED APPLICATIONS

01 Research: Higher Education and Supercomputing
COMPUTATIONAL CHEMISTRY AND BIOLOGY
NUMERICAL ANALYTICS
PHYSICS
06 Defense and Intelligence
06 Computational Finance
07 Manufacturing: CAD and CAE
COMPUTATIONAL FLUID DYNAMICS
COMPUTATIONAL STRUCTURAL MECHANICS
COMPUTER AIDED DESIGN
ELECTRONIC DESIGN AUTOMATION
09 Media and Entertainment
ANIMATION, MODELING AND RENDERING
COLOR CORRECTION AND GRAIN
MANAGEMENT
COMPOSITING, FINISHING AND EFFECTS
EDITING
ENCODING AND DIGITAL DISTRIBUTION
ON-AIR GRAPHICS
ON-SET, REVIEW AND STEREO TOOLS
SIMULATION
WEATHER AND CLIMATE FORECASTING
13 Oil and Gas
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 01
Research: Higher Education and Supercomputing
COMPUTATIONAL CHEMISTRY AND BIOLOGY
Molecular Dynamics
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
ACEMD GPU simulation of molecular mechanics
force fields, implicit and explicit solvent
Written for use only on GPUs Yes
AMBER Suite of programs to simulate molecular
dynamics on biomolecule
PMEMD Explicit Solvent and GB Implicit
Solvent
Yes
CHARMM MD package to simulate molecular
dynamics on biomolecule
Implicit (5x), Explicit (2x) Solvent
via OpenMM. Native CUDA port in
development
Yes
DESMOND High-speed molecular dynamics
simulations of biological systems on
conventional commodity clusters.
The code uses novel parallel algorithms
and numerical techniques to achieve high
performance and accuracy
Yes
DL-POLY Simulate macromolecules, polymers, ionic
systems, etc on a distributed memory
parallel computer
Two-body forces, Link-cell pairs, Ewald
SPME forces, Shake VV
Yes
Folding@Home A distributed computing project that
studies
protein folding, misfolding, aggregation, and
related diseases
Powerful distributed computing molecular
dynamics system; implicit solvent and
folding
Yes
GPUGrid.net A distributed computing project that
uses
GPUs for molecular simulations
High-performance all-atom biomolecular
simulations; explicit solvent and binding
Yes
GROMACS Simulation of biochemical molecules with
complicated bond interactions
Implicit (5x), Explicit (2x) Solvent Yes
HALMD Large-scale simulations of simple and
complex liquids
Simple fluids and binary mixtures (pair
potentials, high-precision NVE and NVT,
dynamic correlations)
Single only
HOOMD-Blue Particle dynamics package written
grounds
up for GPUs
Written for use only on GPUs Yes
LAMMPS Classical molecular dynamics package
Lennard-Jones, Gay-Berne, Tersoff, and
many more potentials
Yes
NAMD Designed for high-performance simulation
of large molecular systems
Full electrostatics with PME and most
simulation features; 100M atom capable
Yes
OpenMM Library and application for molecular
dynamics for HPC with GPUs
Implicit and explicit solvent, custom forces Yes
Test Drive the World's
Fastest Accelerator – Free!
Take the GPU Test Drive, a free and easy way to
experience accelerated computing on GPUs. You
can
run your own application or try one of the preloaded
ones, all running on a remote cluster. Try it today.
www.nvidia.com/gputestdrive
02 | POPULAR GPU-ACCELERATED
APPLICATIONS CATALOG | MAR14
Quantum Chemistry
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Abinit Allows to find total energy, charge density
and electronic structure of systems made of
electrons and nuclei within DFT
Local Hamiltonian, non-local Hamiltonian,
LOBPCG algorithm, diagonalization/
orthogonalization. Contains BigDFT.
Yes
ACES III Takes best features of parallel
implementations of quantum chemistry
methods for electronic structure
Integrating scheduling GPU into SIAL
programming language and SIP runtime
environment
Yes
ADF Density Functional Theory (DFT) software
package that enables first-principles
electronic structure calculations
Fock Matrix, Hessians Yes
BigDFT Implements density functional theory
by solving the Kohn-Sham equations
describing the electrons in a material
DFT; Daubechies wavelets, part of Abinit Yes

CASINO Code for performing Quantum Monte Carlo
(QMC) electronic structure calculations for
finite and periodic systems
TBD Yes
CP2K Program to perform atomistic and
molecular simulations of solid state, liquid,
molecular and biological systems
DBCSR (space matrix multiply library) Yes
GAMESS-UK The general purpose ab initio
molecular
electronic structure program for performing
SCF-, DFT- and MCSCF-gradient
calculations
(ss|ss) type integrals within calculations
using Hartree-Fock ab initio methods
and density functional theory. Supports
organics and inorganics.
Yes
GAMESS-US Computational chemistry suite used to
simulate atomic and molecular electronic
structure
Libqc with Rys Quadrature Algorithm,
Hartree-Fock, MP2 and CCSD
Yes
Gaussian
(In development)
Predicts energies, molecular structures,
and vibrational frequencies of molecular
systems
Joint NVIDIA, PGI and Gaussian
collaboration
Yes
GPAW Real-space grid DFT code written in C and
Python
Electrostatic poisson equation,
orthonormalizing of vectors, residual
minimization method (rmm-diis)
Yes
Jaguar A high-performance ab initio package
for performing gas and solution phase
simulations
TBD Yes
LATTE Density matrix computations CU_BLAS, SP2
Algorithm Yes
MOLCAS Methods for calculating general electronic
structures in molecular systems in both
ground and excited states
CU_BLAS Single only
MOLPRO Used for accurate ab initio quantum
chemistry calculations
Density-fitted MP2 (DF-MP2), density fitted
local correlation methods (DF-RHF, DFKS), DFT
Yes
MOPAC2013 Semiempirical Quantum Chemistry
Pseudodiagonalization, full diagonalization,
and density matrix assembling
Single only
NWChem Calculations Triples part of Reg-CCSD(T),
CCSD and
EOMCCSD task schedulers
Yes
Octopus Used for ab initio virtual experimentation
and quantum chemistry calculations
Full GPU support for ground-state,
real-time calculations; Kohn-Sham
Hamiltonian, orthogonalization, subspace
diagonalization, poisson solver, time
propagation
yes
Q-CHEM Computational chemistry package designed
for HPC clusters
Various features including RI-MP2 TBD
QUICK QUICK is a GPU-enabled ab intio quantum
chemistry software package
Running Hartree-Fock and DFT energy on
GPU, Supports s, p, d, f orbitals on energy
calculation, HF gradient with s,p,d orbital
support, GPU-based ERI generator
Yes
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 03
TeraChem Quantum chemistry software designed to
run on NVIDIA GPU
Full GPU-based solution; Performance
compared to GAMESS CPU version
Yes
Materials Science
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
gWL-LSMS Materials code for investigating the
effects
of temperature on magnetism
Generalized Wang-Landau method Yes
PEtot First principles materials code that
computes the behavior of the electron
structures of materials
Density functional theory (DFT) plane wave
pseudopotential calculations
Yes
QMCPACK Solves the many-body Schrodinger
equation
for electronic structures using a quantum
Monte Carlo method
Main features Yes
Quantum Espresso/
PWscf
An integrated suite of computer codes
for electronic structure calculations and
materials modeling at the nanoscale
PWscf package: linear algebra (matix
multiply), explicit computational kernels,
3D FFTs
Yes
VASP First principles materials code that
computes the behavior of electronic
structures based on quantum theory
Hybrid Hartree-Fock DFT functionals
including exact exchange
Yes
Visualization and Docking Software
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Amira 5 A multifaceted software platform
for visualizing, manipulating, and
understanding Life Science and bio-medical
data
3D visualization of volumetric data and
surfaces
Single only
BINDSURF A virtual screening methodology that
uses

GPUs to determine protein binding sites
Allows fast processing of large ligand
databases
Single only
BUDE Molecular docking program Empirical Free
Energy Forcefield Single only
Core Hopping Rapid screening of novel cores to
improve
drug properties
GPU accelerated application Yes
FastROCS Molecule shape comparison application
Real-time shape similarity searching/
comparison
Yes
Interactive Molecule
Visualizer
Molecular visualization based on a raytracing engine
Written for use only on GPUs Yes
Molegro Virtual Docker 6 Method for performing high
accuracy
flexible molecular docking
Energy grid computation, pose evaluation
and guided differential evolution
Single only
PIPER Protein Docking Protein-protein docking
program Molecule docking TBD
PyMol User-sponsored molecular visualization
system on an open-source foundation
Lines: 460% increase
Cartoons: 1246% increase
Surface: 1746% increase
Spheres: 753% increase
Ribbon: 426% increase
Single only
VEGA ZZ Molecular Modeling Toolkit Virtual logP,
molecular surface values Single only
VMD Visualization and analyzing large biomolecular
systems in 3-D graphics
High quality rendering, large structures
(100M atoms), analysis and visualization
tasks, multiple GPU support for display of
molecular orbitals
Yes
Bioinformatics
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
BarraCUDA Sequence mapping software Alignment
of short sequencing reads,
alignment of indels with gap openings and
extensions
Yes
CUDASW++ Open source software for Smith-
Waterman
protein database searches on GPUs
Parallel search of Smith-Waterman
database
Yes
CUSHAW Parallelized short read aligner Parallel,
accurate long read aligner for
large genomes
Yes
04 | POPULAR GPU-ACCELERATED
APPLICATIONS CATALOG | MAR14
GATK The Genome Analysis Toolkit or GATK
is a software package developed at the
Broad Institute to analyse next-generation
resequencing data. The toolkit offers a wide
variety of tools, with a primary focus on
variant discovery and genotyping as well as
strong emphasis on data quality assurance.
Variant Calling; > 70X speedup Yes
GPU-BLAST Local search with fast k-tuple heuristic
Protein alignment according to BLASTP Single only
mCUDA-MEME Ultrafast scalable motif discovery
algorithm
based on MEME
Scalable motif discovery algorithm based
on MEME
Yes
MUMmerGPU High-throughput local sequence
alignment
program
Aligns multiple query sequences against
reference sequence in parallel
TBD
NVBIO NVBIO is an open source C++ library
of reusable components designed to
accelerate bioinformatics applications using
CUDA.
Data structures, algorithms, and utility
routines useful for building complex
computational genomics applications on
CPU-GPU systems
Yes
NVBowtie A largely complete implementation of the
Bowtie2 aligner on top of NVBIO
Good coverage of Bowtie2 features and
comparable quality results
Yes
REACTA A modified version of GCTA with improved
computational performance, support for
Graphics Processing Units (GPUs), and
additional features. The purpose of REACTA
is to quantify the contribution of genetic
variation to phenotypic variation for complex
traits.
GRM creation, REML analysis, Regional
Heritability (including multi-GPU)
Yes
SeqNFind SeqNFind® is a powerful tool suite that
addresses the need for complete and
accurate alignments of many small
sequences against entire genomes utilizing
a unique hardware/software cluster system
for facilitating bioinformatics research in
Next Generation sequencing and genomic
comparisons.
Hardware and softare for reference
assembly, BLAST, SW HMM, and de novo
assembly
Yes
SOAP3 GPU-based software for aligning short
reads with a reference sequence.It can find
all alignments with k mismatches, where k
is chosen from 0 to 3
Short read alignment tool that is not
heurisitic based, reports all answers
Yes
SOAP3-dp SOAP3-dp: Ultra-fast GPU-based tool for
short read alignment via index-assisted
dynamic programming

Borrows-Wheeler Transformation,
Dynamic Programming, Multi-GPU support
Yes
UGENE Opensource Smith-Waterman for SSE/
CUDA, Suffix array based repeats finder and
dotplot
Fast short read alignment Yes
WideLM Fits numerous linear models to a fixed
design and response
Parallel linear regression on multiple
similarly-shaped models
Yes
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 05
NUMERICAL ANALYTICS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
ArrayFire Comprehensive GPU function library
Hundreds of functions for math, signal/
image processing, statistics, and more.
Yes
Mathematica Wolfram A symbolic technical
computing language
and development environment
Development environment for CUDA and
OpenCL
Yes
MATLAB by Mathworks GPU acceleration for
MATLAB (high-level
technical computing language)
Support for 200+ of most used MATLAB
functions (incl. Signal Processing, Image
Processing, Communications Systems, etc)
Yes
NMath Premium GPU-accelerated math and
statistics for
.NET, automatically detects the presence
of a CUDA-enabled GPU at runtime
and seamlessly redirects appropriate
computations to it.
Automatically offloads computations to the
GPU
Single only
PHYSICS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Chroma Lattice Quantum Chromodynamics (LQCD)
Wilson-clover fermions, Krylov solvers,
Domain-decomposition
Yes
ENZO 3D block-structured AMR code for
cosmological structure formation
Accelerated magneto hydrodynamics
solvers
Yes
GTC Simulates microturbulence and transport in
magnetically confined fusion plasma
Electron push and shift (accounting for
>80% of run time)
Yes
GTS Simulates microturbulence and the motion
of charged particles and interactions in
fusion plasma
Push and shift for both electron and ion
dynamics
Yes
MILC Lattice Quantum Chromodynamics (LQCD)
codes simulate how elemental particles
are formed and bound by the "strong force"
to create larger particles like protons and
neutrons
Staggered fermions, Krylov solvers,
Gauge-link fattening
Yes
PIConGPU A relativistic Particle-in-Cell code that
describes the dynamics of a plasma by
computing the motion of electrons and ions
subject to the Maxwell-Vlasov equation.
Simulation of laser-wakefield acceleration
of electrons
Yes
QUDA Library for Lattice QCD calculations using
GPUs
CUDA supports the following fermion
formulations: Wilson,Wilson-clover,Twisted
mass,Improved staggered (asqtad or HISQ)
and Domain wall
Yes
RAMSES Simulates astrophysical problems on
different scales (e.g. star formation,
galaxy dynamics, cosmological structure
formation)
CUDA acceleration is applied for radiative
transfer for reionization, and the
hydrodynamic solver using AMR
Yes
Chemora Chemora is a system for performing
simulations of systems described
by differential equations running on
accelerated computational clusters
Chemora embeds the equations'
computational kernels into dynamically
compiled loop nests shaped for input size
and GPU structure
Yes
06 | POPULAR GPU-ACCELERATED
APPLICATIONS CATALOG | MAR14
Defense and Intelligence
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
DigitalGlobe Advanced
Ortho Series
Geospatial Visualization Image orthorectification Yes
Eternix Blaze Terra Geospatial Visualization 3D
visualization of geospatial data Yes
Exelis (ITT) ENVI Geospatial Visualization Image
orthorectification
(custom builds only)
Yes
MrGeo Geospatial Visualization Terrain analytics Yes
GeoWeb3d Desktop Geospatial Visualization 3D
visualization of geospatial data Yes
Incogna GIS Geospatial Visualization Image
processing on Tesla cloud servers
Object recognition
Yes
Intergraph Motion Video
Analyst
Video filters and mosaic'ing - Geo-fuses
FMV analytics with intelligence data
Full motion video ortho mosaic processing Single

only
Intuvision Panoptes 3.0 Video Analytics Object
recognition and change detection Yes
LuciadLightspeed Geospatial Visualization and
Analysis Geospatial Situational Awareness Single
only
MotionDSP Ikena ISR Real-time Full Motion Video
and WAMI
enhancement and analytics
Real-time super-resolution-based video
enhancement, filtering, mosaicing, video
analytics, and transcoding
Yes
Nerve ViD SrX Video Analytics Object recognition
and tracking Yes
NerVve Visual Search
Solution (NVSS)
Video/Image Live and Forensic Search Video and
Image content search Yes
OpCoast SNEAK Electromagnetic signals
propagation
modeling for complex urban and terrain
environments
Ray tracing, DTED and remote sensing
inputs
Yes
PCI Geomatics GXL Geospatial Visualization Image
orthorectification and additional
image processing
Yes
GAIA Multi-GPU, Multi-Machine distributed
object store providing SQL style query
capability, advanced geospatial query
capability,heatmap generation, and
distributed rasterization services
Built for GPUs, scales to many machines
and many cards without end user query
modification or pre-meditation
Yes
Computational Finance
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Aaon Benfield
Pathwise™
Specialized platform for real-time hedging,
valuation, pricing and risk management
Spreadsheet-like modeling interfaces,
Python-based scripting environment and
Grid middleware
Yes
Hanweck Associates Real-time options analytical
engine (Volera) Real-time options analytics engine
Yes
Murex MACS Analytics
Library
Analytics library for modeling valuation and
risk for derivatives across multiple asset
classes.
Market standard models for all asset
classes paired with the most efficient
resolution methods (Monte Carlo
simulations and Partial Differential
Equations)
Yes
Numerical Algorithms
Group (NAG)
Random number generators, Brownian
bridges, and PDE solvers
Monte Carlo and PDE solvers Single only
RMS Catastrophic risk modeling for FSI
(earthquakes, hurricanes, terrorism,
infectuous diseases)
Risk analytics Yes
Tanay ZX Lib (Fuzzy
Logic)
Financial analytics and data mining library Monte
Carlo simulations, pricing of vanilla
and exotic options, fixed income analytics,
data mining
Yes
SciComp, Inc Derivative pricing (SciFinance) Monte
Carlo and PDE pricing models Single only
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 07
Xcelerit SDK Software Development Kit (SDK) to
boost
the performance of Financial applications
(e.g. Monte-Carlo, Finite-difference) with
minimum changes to existing code
C++ programming language, crossplatform (back-end
generates CUDA and
optimized CPU code), supports Windows
and Linux operating systems.
Yes
Synerscope's Synerscope
Data Visualization
Visual big data exploration and insight tools Graphical
exploration of large network
datasets including geo-spatial and
temporal components
Single only
QuantAlea's Alea.cuBase
F#
F# package enabling a growing set of F#
capability to run on a GPU.
F# for GPU accelerators Yes
Altimesh's Hybridizer C# Multi-target C# framework
for data parallel
computing.
C# with translation to GPU or Multi-Core
Xeon
Yes
MISYS Global Risk Regulatory compliance and
enterprise wide
risk transparency package
Risk Analytics Yes
MiAccLib High Speed Multi-Algorithm Search Engine
library providing high speed text string
search with scalability of searching text
and/or keywords on hundreds millions of
records and/or text data.
Exact Text match Search, Approximate\
Similarity Text Search, Wild Card
Text Search, Proximity & Percentage
Text Search, MultiKeyword and
MultiColumnMultiKeyword Text Search,
RadixSort Text
Yes
SunGard Adaptiv
Analytics
A flexible and extensible engine for fast

calculations of a wide variety of pricing and
risk measures on a broad range of asset
classes and derivatives
Existing models code in C# supported
transparently, with minimal code changes,
Supports multiple backends including
CUDA and OpenCL, Switches transparently
between multiple GPUs and CPUS
depending on the deal support and load
factors
Yes
Manufacturing: CAD and CAE
COMPUTATIONAL FLUID DYNAMICS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Altair AcuSolve General purpose CFD software
Linear equation solver Yes
ANSYS Fluent General purpose CFD software
Radiation heat transfer model, linear
equation solver
Yes
Autodesk Moldflow Plastic mold injection software
Linear equation solver Single only
CPFD Barracuda-VR and
Barracuda
Fluidized bed modeling software Linear equation
solver, particle
calculations
Single only
FluiDyna Culises for
OpenFOAM
Solver library for general purpose CFD
software
Linear equation solvers Yes
FluiDyna LBultra General purpose CFD software
Lattice-Boltzmann solver Yes
Prometech Particleworks Particle-based CFD
software Implicit and explicit solvers Single only
Vratis ARAEL General purpose CFD software based
on
FVM with OpenFOAM compatibility
Linear equation solver Yes
Turbostream Ltd. CFD software for turbomachinery
flows Explicit solver Yes
Vratis SpeedIT extreme
for OpenFOAM
Solver library for general purpose CFD
software
Linear equation solvers Yes
RESEARCH CFD DEVELOPMENTS
FEFLO (GMU-Lohner) General purpose CFD
software for
compressible and incompressible flows
Implicit and explicit solver Yes
SD++ (StanfordJameson)
General purpose CFD software for
compressible flows.
Explicit solver Yes
S3D (Sandia and Oak
Ridge NL)
Direct numerical solver (DNS) for turbulent
combustion
Chemistry model Yes
DualSPHysics SPH-based CFD software SPH model
Yes
NASA FUN3D General purpose CFD software Linear
equation solver Single only
08 | POPULAR GPU-ACCELERATED
APPLICATIONS CATALOG | MAR14
COMPUTATIONAL STRUCTURAL MECHANICS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Abaqus/Standard Simulation and analysis tool for
structural
mechanics
Direct sparse solver Yes
ANSYS Mechanical Simulation and analysis tool for
structural
mechanics
Direct and iterative solvers Yes
Impetus Afea Predicts large deformations of
structures
and components exposed to extreme
loading conditions
Linear equation solver Yes
LS-DYNA Implicit Multiphysics simulation package
used Linear equation solver Yes
MSC Nastran Simulation and analysis tool for
structural
mechanics
Direct sparse solver Yes
MSC Marc Simulation and analysis tool for structural
mechanics
Direct sparse solver Yes
OptiStruct Simulation and analysis tool for structural
mechanics
Direct and iterative solvers Yes
DS Exsight Uses Abaqus Standard for GPU
computing Direct sparse solver Single only
DS DesignSight Uses Abaqus Standard for GPU
computing Direct sparse solver Single only
PAM-CRASH Implicit Multiphysics simulation
package used Linear equation solver Single only
NX Nastran Simulation and analysis tool for structural
mechanics
Linear equation solver Single only
COMPUTER AIDED DESIGN
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
AutoCAD 2D and 3D CAD design, drafting, modeling,
architectural drawing, and engineering
software. Supports Open GL. Native DWG™
support.
Surface, mesh, and solid modeling tools,
model documentation tools, parametric
drawing capabilities. Native DWG™
support.
Single only
AutoCAD Design Suite AutoCAD 2014 software, plus
tools to create,
capture, connect, and showcase designs
2D/3D display of designs, interactive 3D
presentation with realistic materials,
rendering-ray tracing
Single only
Autodesk 3ds Max 3D animation creative toolset for
modeling,
animation, simulation, and rendering for
games, film, and motion graphics
3D modeling, mesh and surface modeling,
improved Nitrous viewport performance,

iray rendering
Yes
Inventor 3D mechanical design, documentation, and
product simulation
Uses BIM for intelligent building
components to improve design accuracy
Single only
Revit Building Information Modeling (BIM) for
structural engineering, and construction
Modeling (BIM) to design, build, and
maintain higher-quality, more energyefficient
buildings
Single only
NVIDIA Iray A ready-to-integrate, physically-based,
photorealistic rendering solution
Iray Interactive; Iray Photoreal; Iray
Cluster. Fast interactive ray tracing;
Physically-based, global-illumination
rendering; Distributed cluster rendering
Yes
Dassault Systemes CATIA
- Live Rendering
Photorealistic rendering Interactive. Fully integrated
in CATIA V6.
Network rendering
Yes
Bunkspeed Suite Easy to use photorealistic rendering
software
Iray-based ray-tracing. Animation support.
Network rendering
Yes
RTT DeltaGen Redefines high-end 3D visualization
and
realtime interaction. This latest version
gives users a broad suite of robust new
features to truly revolutionize processes
and help increase visual quality, speed, and
flexibility.
Interactive ray tracing and global
illumination. Integration with Siemens
TeamCenter. Cluster support Realtime
& Offline Production Process Integration
and scene building. Scene Analysis, Xplore
DeltaGen, SDK for DeltaGen
Yes
PTC Creo Parametric Parametric design solution
suite Anti-aliasing, better lighting and enhanced
shaded-with-edges mode. GRID Support
Single only
Siemens PLM NX and
Teamcenter
Product lifecycle management solutions
from design to simulation to production to
service
Design software, NX, and PLM viewer
applications, TcVis and Active Workspace
Client.
Single only
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 09
ELECTRONIC DESIGN AUTOMATION
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Agilent Technologies ADS Simulation tool for design
of RF, microwave
and high speed digital circuits
Signal integrity simulation Single only
Agilent Technologies
EMPro
Modeling and simulation environment for
analyzing 3D EM effects of high speed and
RF/Microwave components
FDTD solver Yes
ANSYS Nexxim Circuit simulation engine for
RF/analog/
mixed-signal IC design; IBIS-AMI analysis
speedup with GPU computing.
AMI analysis Single only
ANSYS HFSS Simulation tool for modeling 3-D full-
wave
electromagnetic fields in high-frequency
and high-speed electronic components.
Transient solver Yes
CST Microwave Studio
(MWS)
3D EM modeling and simulation Transient solver
Yes
Delcross Savant Simulation tool for installed antenna
performance and antenna-to-antenna
coupling
High-frequency solver Yes
EMSS FEKO 3D EM modeling and simulation MoM
solver Single only
Remcom XFdtd 3D EM modeling and simulation
FDTD solver Yes
SPEAG SEMCAD-X 3D EM modeling and simulation
FDTD solver Yes
Rocketick RocketSim Verilog simulation Verilog
simulation Yes
Media and Entertainment
ANIMATION, MODELING AND RENDERING
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Autodesk 3ds Max +
NVIDIA iray
3D modeling, animation, and rendering iray
interactive, photorealistic and
physically correct rendering
Yes
Autodesk Maya 3D modeling, animation, and
rendering Increased model complexity, larger scenes,
windowed stereo
Single only
Autodesk Motion Builder Character animation and
motion capture Increased model complexity at
interactive
rates
Single only
Autodesk Mudbox 3D sculpting Increased model
complexity at interactive
rates
Single only
Cebas finalRender GPU Renderer CUDA-based GPU
Rendering Yes
CentiLeo GPU Render GPU Renderer CUDA-based
GPU Rendering Single only
Chaos V-Ray RT GPU Renderer CUDA interactive
and final-frame GPU
Rendering
Yes
Jawset TurbulenceFD Physics-based simulation plug-

in Maximus supported GPU simulation using
CUDA
Single only
Maxon Cinema 4D 3D modeling, animation, and
rendering Increased model complexity at interactive
rates
Single only
NewTek Lightwave 3D modeling, animation, and
rendering Increased model complexity at interactive
rates
Single only
Otoy Octane Render GPU Renderer CUDA-based
GPU final-frame rendering Yes
Pixologic Sculptris 3D sculpting Increased model
complexity at interactive
rates
Single only
Redshift Renderer GPU-accelerated, biased renderer
CUDA-based GPU final-frame rendering Yes
Side Effects Houdini 3D modeling, animation, and
rendering Maximus-supported GPU simulation using
OpenCL
Single only
The Foundry Mari 3D paint Increased model
complexity at interactive
rates
Single only
10 | POPULAR GPU-ACCELERATED
APPLICATIONS CATALOG | MAR14
COLOR CORRECTION AND GRAIN
MANAGEMENT
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Adobe SpeedGrade CC Color grading Real-time
grading and finishing with
Lumetri Deep Color Engine
Single only
Assimilate Scratch Color grading and finishing
Accelerated debayering for real-time
digital finishing
Single only
Blackmagic DaVinci
Resolve
Color grading Real-time color correction, real-time
denoising
Yes
Cinnafilm Dark Energy Application and plug-in for
image
enhancement
Image de-noising and restoration.
Simultaneous active and background
rendering
Yes
Digital Vision Nucoda Color grading Real-time color
correction Single only
Marquise Technologies
Rain
Color grading CUDA-based real-time color correction
Single only
Quantel Pablo Rio Color grading and finishing Real
time color correction Yes
Red Digital Cinema
REDCINE-X
Color grading CUDA-accelerated debayering Single
only
SGO Mistika Color grading and finishing Real-time
color correction and finishing Single only
The Pixel Farm PFClean Image restoration and
remastering CUDA-based image processing
acceleration
Single only
COMPOSITING, FINISHING AND EFFECTS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
ABSoft Neat Video Video noise reduction plug-in
CUDA-based acceleration Single only
Adobe After Effects CC Motion graphics and effects
3D raytracing engine based on NVIDIA
OptiX
Yes
Autodesk Flame
Premium
Finishing and color grading Post-production
integrated end-to-end
toolset for 3D VFX, editorial, and color
grading
Single Only
Autodesk Smoke Finishing and editing Faster effects
Single only
Boris FX Continuum
Complete
Visual effects plug-in Faster effects Single only
Cinnafilm Dark Energy
Plug-in
Color management All features accelerated with
CUDA Yes
CoreMelt complete Visual effects plug-in Faster
effects Single only
eyeon Fusion Effects and compositing Faster effects
Single only
GenArts Monsters GT Visual effects plug-in Faster
effects Single only
GenArts Sapphire Visual effects plug-in Faster
effects Single only
ROBUSKEY Chroma keyer plug-in Faster effects
Single only
Neat Video Open FX Video noise reduction plug-in
Faster effects Single only
NewBlueFX Video
Essentials
Video effects plug-in Faster effects Single only
NewBlue Titler Pro Video titling plug-in Faster effects
Single only
Pixelan AnyFX Video effects plug-in Faster effects
Single only
Re:Vision Effects Visual effects plug-in Faster effects
Single only
Red Giant Effects Suite Visual effects plug-ins Faster
effects Single only
Red Giant Magic Bullet
Looks
Color and finishing tools Faster effects Single only
The Foundry HIERO Shot management, conform
and review
timeline
Better interactivity Single only
The Foundry NUKE and
NUKEX
Compositing tools with 3D tracker Faster effects
Single only
Video Copilot Software

Element 3D
3D object based particle system Faster effects Yes
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 11
Video Copilot Optical
Flares
Lens flares plug-in for After Effects Faster effects
Single only
Video Copilot Twitch Video effects plug-in for After
Effects Faster effects Single only
EDITING
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Adobe Photoshop CC Image editing Over 30 effects
for smoother image
manipulation in Mercury Graphics Engine
Single only
Adobe Premiere Pro CC Video editing Mercury
Playback Engine for real-time
video editing & accelerated rendering
Yes
Apple Final Cut Pro Video editing Faster effects
Single only
Avid Media Composer Video editing Faster video
effects, unique stereo 3D
capabilities
Single only
Grass Valley Edius Video editing Faster effects
Single only
Harris Velocity Video editing Faster effects Single
only
Quantel Qube Broadcast video editing Faster video
effects, unique stereo 3D
capabilities
Single only
Sony Vegas Pro Video editing Faster video effects
and encoding Single only
ENCODING AND DIGITAL DISTRIBUTION
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Cinnafilm Tachyon Standards conversion Video
procession and encoding Yes
Digimetrics Aurora Automated video and audio test
and
measurment
Video encoding and video processing Single only
Elemental Live Live streaming video processing and
encoding
Video encoding and video processing Yes
Elemental Server File-based video processing and
encoding Video encoding and video processing Yes
isovideo Viarte Video standards conversion CUDA-
accelerated video procession and
encoding
Yes
MainConcept CUDA
H.264/AVC Encoder SDK
H.264 video encoder Video encoding and video
processing Yes
Sorenson Squeeze Video transcoding application and
plug-In Video encoding and video processing Yes
Snell Alchemist on
Demand
Video standards conversion Video procession and
encoding Yes
Telestream Vantage Video transcoding and
processing Video encoding and video processing Yes
ON-AIR GRAPHICS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Avid On-air Graphics Motion Graphics Real-time
graphics rendering Single only
Brainstorm Estudio Virtual Sets and motion graphics
Real-time graphics rendering Single only
Chyron HyperX On-air graphics Real-time graphics
rendering Single only
Chyron LEX On-air graphics Real-time graphics
rendering Single only
Harris Inscriber On-air graphics Real-time graphics
rendering Single only
Monarch dScript 3D Virtual sets Real-time graphics
rendering Single only
Monarch Twister Pro Virtual Sets and motion
graphics Real-time graphics rendering Single only
Monarch Virtuoso Virtual Sets and motion graphics
Real-time graphics rendering Single only
Pixel Power Clarity On-air graphics Real-time
graphics rendering Single only
RT Software tOG On-air graphics Real-time graphics
rendering Single only
Vizrt Viz Engine On-air graphics Real-time graphics
rendering Single only
Wasp 3D Beehive On-air graphics and virtual sets
Real-time graphics rendering Single only
12 | POPULAR GPU-ACCELERATED
APPLICATIONS CATALOG | MAR14
ON-SET, REVIEW AND STEREO TOOLS
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
3ality Intellicam 3D stereo camera adjustment CUDA-
based imaging Single only
BlueFish 4K Review Review and approval of 4K
content Real-time Single only
Colorfront On-Set Dailies Review and approval of 4K
content Real-time Yes
eyeon Dimension 3D stereoscopic workflow Real-
time Single only
Lightcraft Previzion On-set workflow Real-time Single
only
Tweak RV Review and approval of 4K content Real-
time Single only
SIMULATION
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
3DAliens Glu3d SPH fluid simulation Faster
simulation Single only
Blastcode Kilton/
Megaton
Physics-based simulation plug in Faster simulation
Single only
Jawset TurbulenceFD Physics-based simulation plug
in Faster simulation Single only,
multi-GPU
available in
late Q2/Q3
WEATHER AND CLIMATE FORECASTING
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Accuweather Cinemative
HD
Weather graphics Real-time Single only

Accuweather Storyteller Weather graphics Real-time
Single only
ASUCA Regional atmospheric model Entire model
Yes
CAM - SE Global atmosphere model for climate
research
Dynamical core Yes
COSMO Regional atmospheric model Entire model
Yes
GEOS-5 Global climate model Entire model Yes
HOMME Dynamical core for global atmospheric
model
Dynamical core Yes
HYCOM Ocean circulation model Dynamical core
Single only
Meteo Earth Weather graphics Real-time Single only
MITgcm Ocean circulation model Dynamical core
Single only
NEMO Ocean circulation model Entire model (GYRE
config) Yes
NIM Dynamical core for global atmospheric
model
Dynamical core Yes
Weather Central Fusion
Studio
Weather graphics Real-time Single only
WRF Regional numerical weather prediction
model
WSM5, WSM3, Ice Microphysics models Yes
WSI TrueView Max Weather graphics Real-time
Single only
POPULAR GPU-ACCELERATED APPLICATIONS
CATALOG | MAR14 | 13
Oil and Gas
APPLICATION DESCRIPTION SUPPORTED
FEATURES MULTI-GPU SUPPORT
Acceleware
AxRTM
AxKTM
Seismic Processing RTM, Kirchhoff, control source,
electromagnetism, forward modeling
Yes
CGGV GeoVation Seismic Processing Multiple
algorithms (RTM, etc) Yes
ffA Geoteric Seismic Interpretation Attributes
calculations, geobodies
extraction
Yes
ffA SVI Pro Seismic Interpretation Attributes
calculations, geobodies
extraction
Yes
ffA SEA3D Pro Seismic Interpretation Attributes
calculations, geobodies
extraction
Yes
GeoStar Seismic Suite Seismic Processing Multiple
algorithms (RTM, etc) Yes
Headwave Suite Seismic Interpretation Attributes
calculations, Volume Rendering Yes
HUE HUEspace Seismic Interpretation Interpretation
development platform Yes
OpenGeo Solutions
OpenSeis
Seismic Processing Spectral Decomposition Yes
Panorama Tech Seismic Processing, Modeling
Multiple algorithms (RTM, etc) Yes
Paradigm Echos RTM Seismic Processing RTM
algorithm Yes
Paradigm SKUA Reservoir Modeling Faults, Horizons
and Flow Simulation Grid Yes
Paradigm Geophysical
VoxelGeo
Seismic Interpretation Volume Rendering, Horizon
Flattening Yes
Roxar RMS Reservoir Modeling Multi GPU
capabilities via HUEspace Yes
Seismic City Prestack
Interpretation
Seismic Processing Multiple algorithms (RTM, etc)
Yes
SpectraSeis Seismic Processing Full elastic wave-
equation imaging and
analysis of microseismic fracture data
Yes
Stoneridge Technologies
GAMPACK
Reservoir Simulation GPU Algebraic MultiGrid
Package Yes
Terraspark Insight Earth Seismic Interpretation
Horizon orientation attributes; automated
fault extraction, Curvature Attributes
Yes
Tsunami RTM Seismic Processing RTM algorithm
Yes
WesterGeco Omega2
RTM
Seismic Processing Multiple algorithms (RTM, etc)
Yes
For more information on GPU-accelerated
applications please visit,www.nvidia.com/teslaapps

Software Development Ecosystem for Intel Xeon Phi

	Open Source	Commercial
Compiler	gcc (kernel build only, not for applications), Python	Intel Parallel Studio XE (C++ & Fortran) Intel Cluster Studio XE (C++ & Fortran) CAPS HMPP compiler (beta)
Debugger	gdb	Intel Debugger Rogue Wave TotalView (beta) Allinea DDT
Libraries	TBB (in Intel Studio XE) MPICH2 FFTW NetCDF	Intel MKL, Intel MPI, OpenMP, Intel IPP, Cilk™ Plus (in Intel Studio XE products), NAG Rogue Wave IMSL
Profiling & Analysis Tools		Intel VTune Amplifier XE Intel Trace Analyzer & Collector Intel Inspector XE
Workload Scheduler		Altair PBS Professional, Adaptive Computing Moab
System Management		SGI Management Center Bright Cluster Manager (beta)

For more information:
http://software.intel.com/en-us/articles/intel-and-third-party-tools-and-libraries-available-with-support-for-intelr-xeon-phitm

Next generation XEON: Haswell

Tick/Tock Development Model

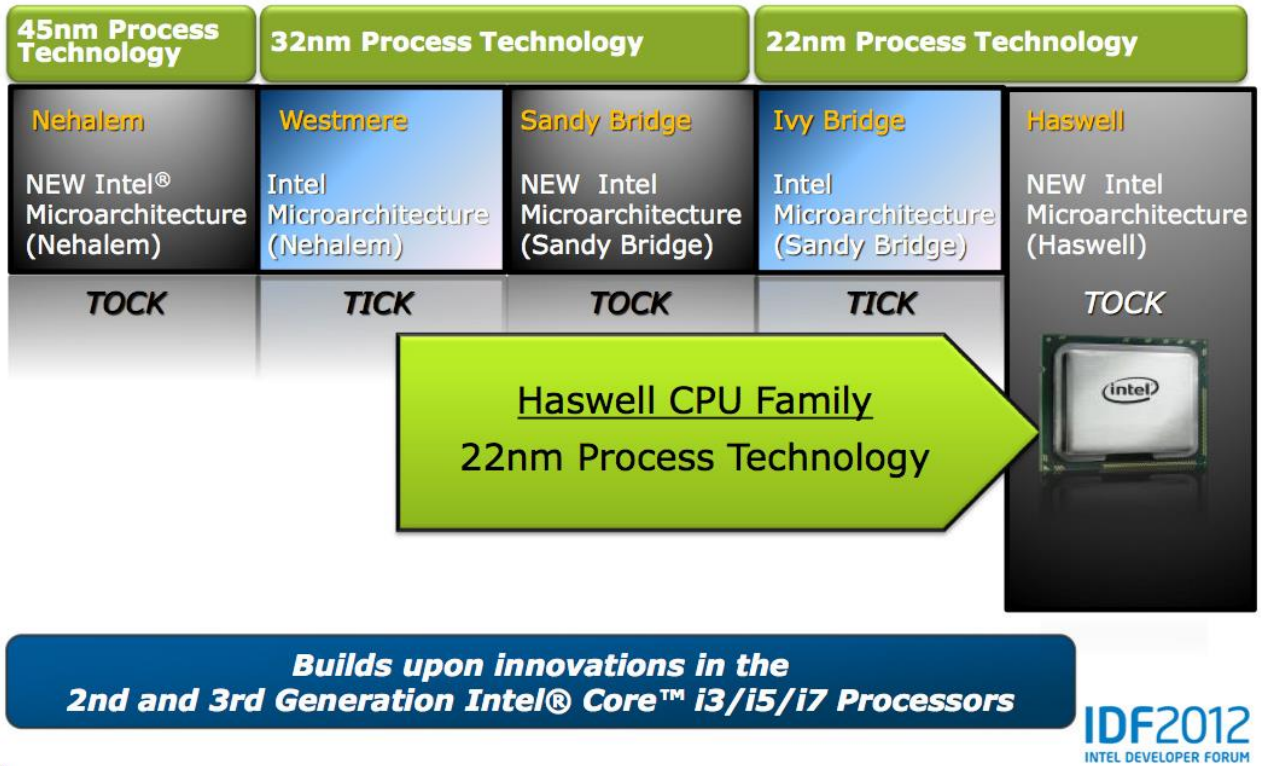

Next-next generation HPC XEON

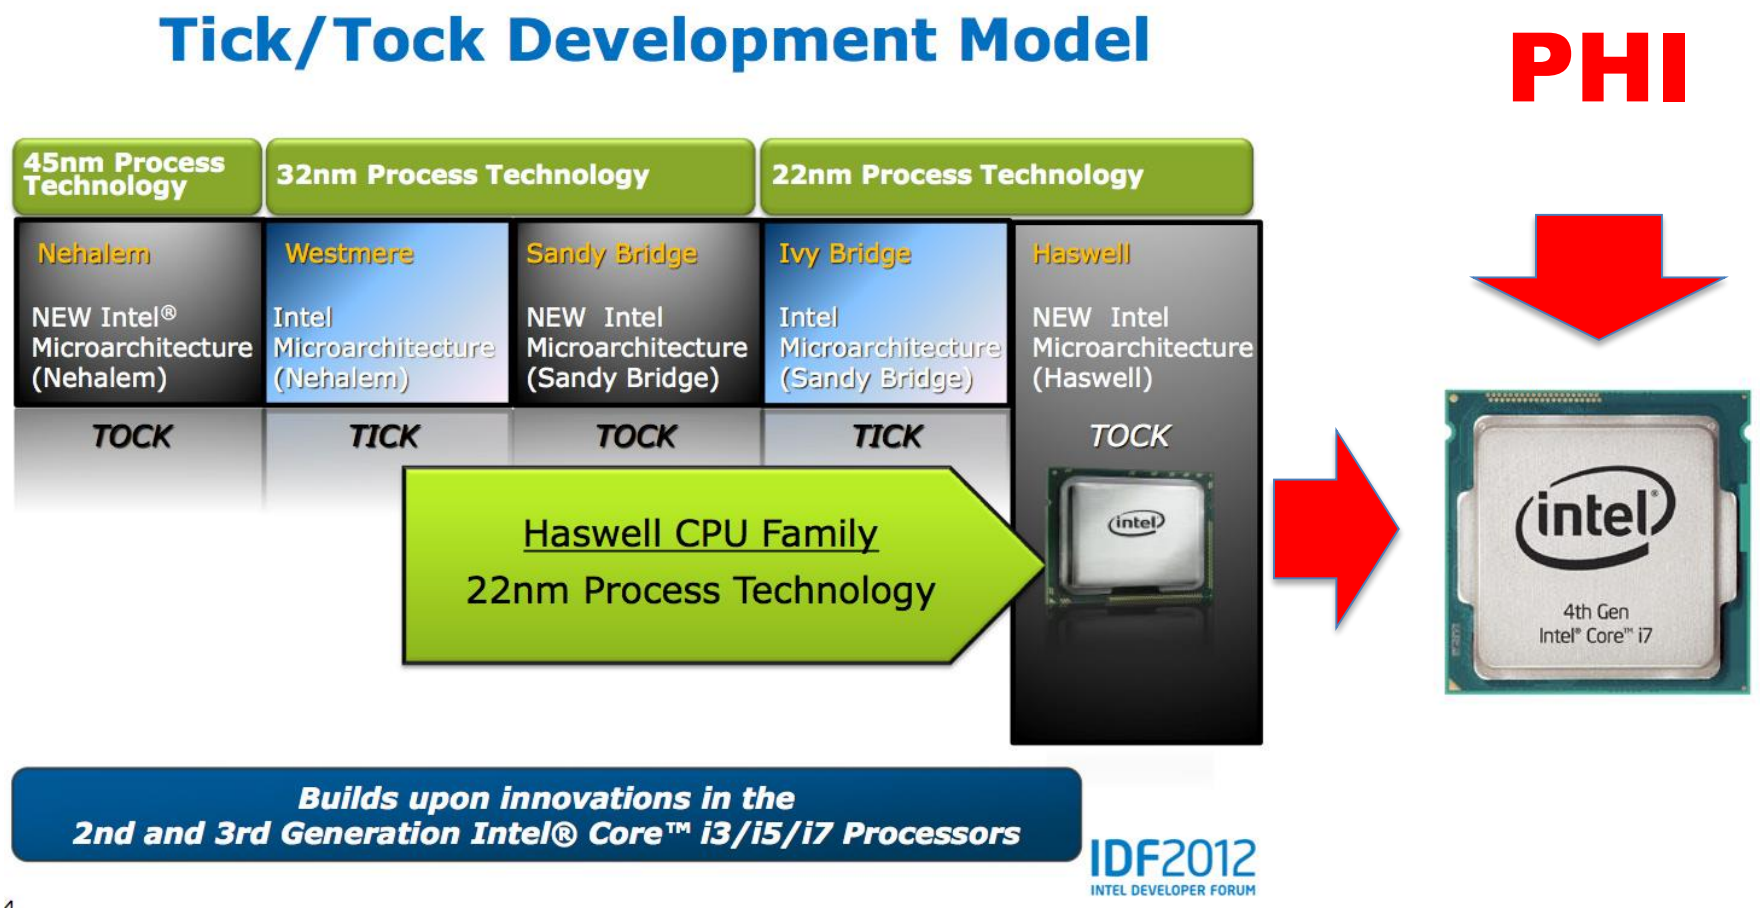

“Pyramid 1U”

NVIDIA GPU accelerator support:
- 3x K10, K20, K20X

NVIDIA graphics GPU support:
- Quadro Plex 7000

Intel accelerator support:
- 3x Phi

Model	C1104G-RP5
Chassis Profile	1U standard-depth
Servers/System	One dual-socket
Chipset	Intel® C600
Max. Processors	Two Intel® Xeon® processor E5-2600 series
Max. CPU TDP	115W
Memory Slots	8 DIMM slots
Memory Type	1600/1333/1066/800 MHz DDR3 ECC Reg
Max. Hard Disk Drives	4 x 2.5" drives
Expansion Slot	**Three PCIe 3.0 x16 double-width slots** One external PCIe 3.0 x8 low-profile slot
Networking (Onboard)	Dual-Port GigE controller (Intel® I350)
IPMI Remote Management	Integrated IPMI 2.0
Power Supply	**1800W** Redundant* Platinum Level

10

“Pyramid 2U”

NVIDIA GPU accelerator support:
- 4x K10, K20, K20X

NVIDIA graphics GPU support:
- Quadro Plex 7000

Intel accelerator support:
- 4x Phi

Model	C2110G-RP5
Chassis Profile	2U standard-depth
Servers/System	One dual-socket
Chipset	Intel® C600
Max. Processors	Two Intel® Xeon® processor E5-2600 series
Max. CPU TDP	130W
Memory Slots	8 DIMM slots
Memory Type	1600/1333/1066/800 MHz DDR3 ECC Reg
Max. Hard Disk Drives	10 x 2.5” drives
Expansion Slot	**Four internal PCIe 3.0 x16 double-width** **(optional riser card for 2 additional external PCIe 3.0 x16 double width slots)**, One external PCIe 3.0 x8 low profile and one PCIe 2.0 x4 full height
Networking (Onboard)	Dual-Port GigE controller (Intel® I350)
IPMI Remote Management	Integrated IPMI 2.0
Power Supply	**1800W** Redundant** Platinum Level

11

SGI UV 2000

Within a Single Standard 19” Rack

Up to:

- 64 socket/512 cores/1024 threads
  - Or **34 CPU + 30 Intel® Xeon® Phi™**
  - Or **36 CPU + 28 NVIDIA® Tesla® K20,K20X, or K40 (2 partitions)**
- 16TB memory
- 63 x16 PCIe Gen 3 Links

12

SGI UV and Scale-out Compared

Feature	Standard Scale-out Servers	SGI UV
Architecture Reference Terms	Scale Out Cluster Distributed Memory	Scale Up Single System Image Shared Memory
System Limit	16 cores, 0.5 TB memory	2048 cores, 64 TB memory
CPU	x86 Intel® Xeon® or AMD Opteron™	X86 Intel® Xeon®
Memory, Storage, networking	Industry standard	Industry Standard
Interconnect or System Fabric	Ethernet or InfiniBand	SGI NUMAlink
Hardware Package or Building Block	Blades or Rackmount, 19” rack	Blades (UV2000) or Rackmount(UV20), 19” rack
Software	Off the Shelf	Off the Shelf
Applications	Small scale single apps cluster or “MPI” apps	Small ***or large*** single apps Cluster or “MPI” apps also
Cost	Lowest cost x86 architecture	20-75% > scale-out, ~1/3 the cost of ‘Big Iron’.

13

SGI® Value Propositions

- Long history of working with accelerators
  - “Home-brewed” – Geometry Engine, TPU (Tensor Processor Unit)
  - FPGA’s (RASC™ technology), GPUs, SOCs
  - 50 application experts focused on it
- Accelerators in both a scale-up and scale-out environment
  - 32 Intel® Xeon® Phi™ coprocessors in SGI® UV™ 2000 (COSMOS and TGAC)
  - SGI: the only vendor able to deliver hybrid scale-out and scale-up solutions
- Everything you need in a powerful solution
  - Factory-integrated, tested, rack level delivery – plug in and go
  - Starter kits
  - Worldwide customer support
- Fully-managed, with SGI Management Center and Performance Suite software

14

SGI® Integrated Clusters for Intel Xeon Phi

- A complete, managed GPU solution of software and hardware, all you need for a powerful deployme
- Hardware stack
  - Rackable™ C2108 head node
  - Either/or Rackable C1104G or C2110G compute nodes
  - Infiniband NICs and Switch
  - Ethernet switch for management
  - Add SGI InfiniteStorage solutions, either direct-attached to the head NAS
- Software stack
  - SGI Management Center software
  - Altair PBSpro Load Management software
- Complete Factory Test & Integration

SGI® Management Center

Premium Edition offers Accelerator Monitoring and Management!

- Single system management console with remote server monitoring and control
- Ease of use, full system management including GUI and CLI
- Policy driven, fine grained po computing
- Advanced fault, event, improved reliability
- Advanced capabilities inclu management, and high avail

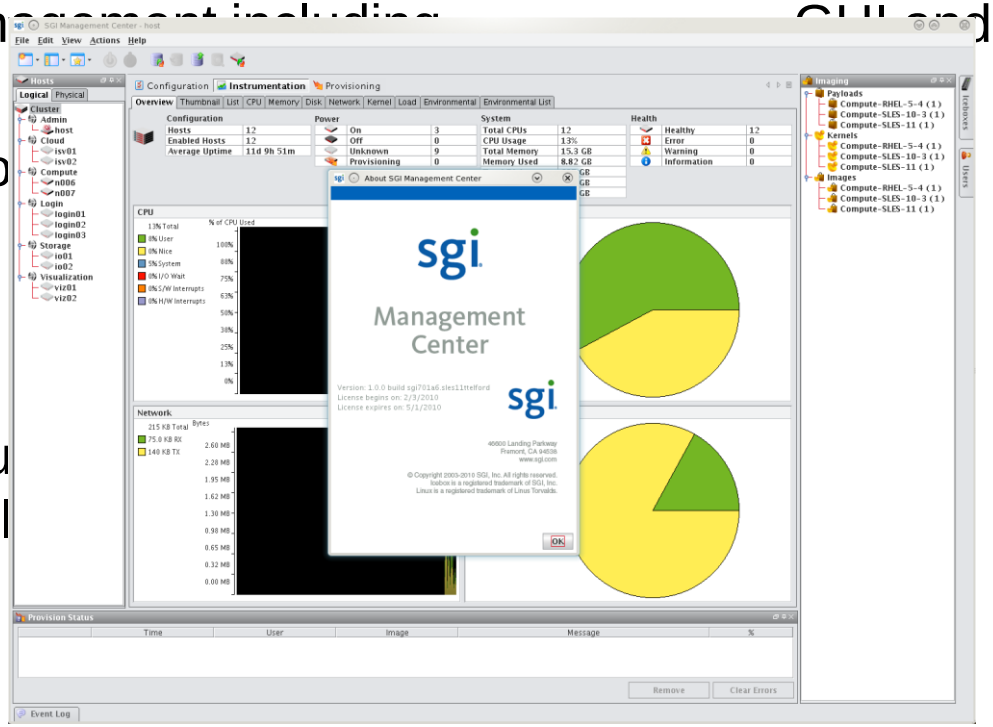

SMC Premium Edition– Accelerator Monitoring

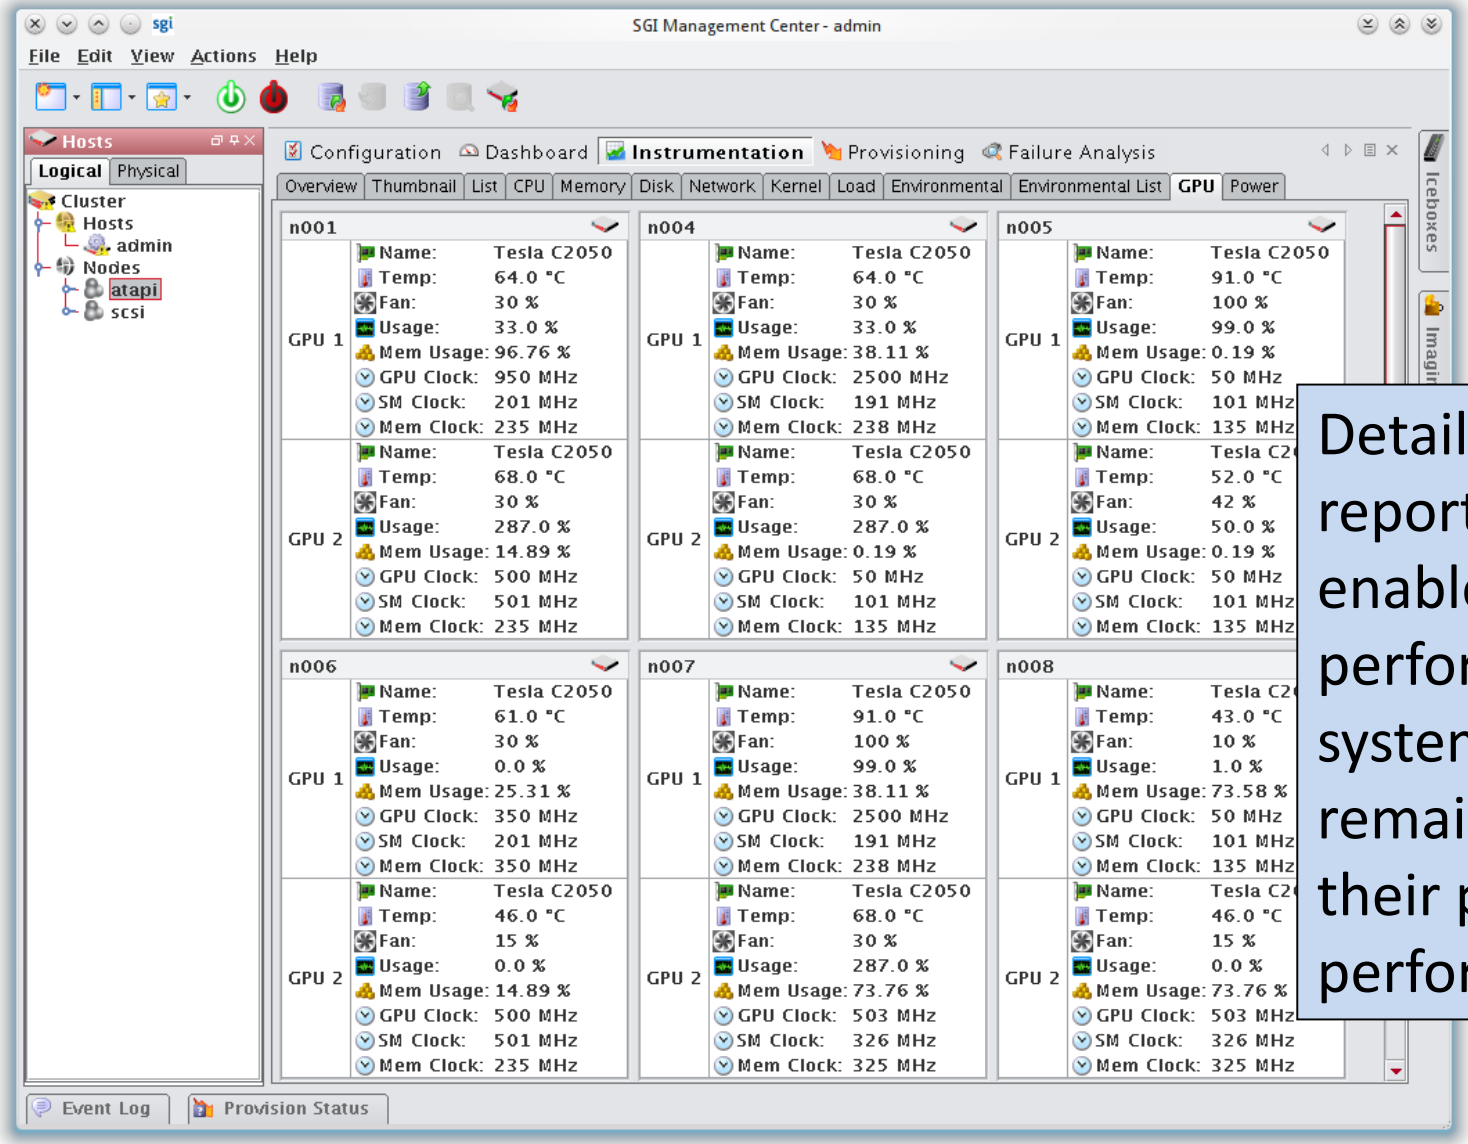

Detailed GPU reporting enables high performance systems to remain at their peak performance

Optimized for HPC & Big Data Solutions

- SGI® ICE™ X
- SGI UV™ 2000

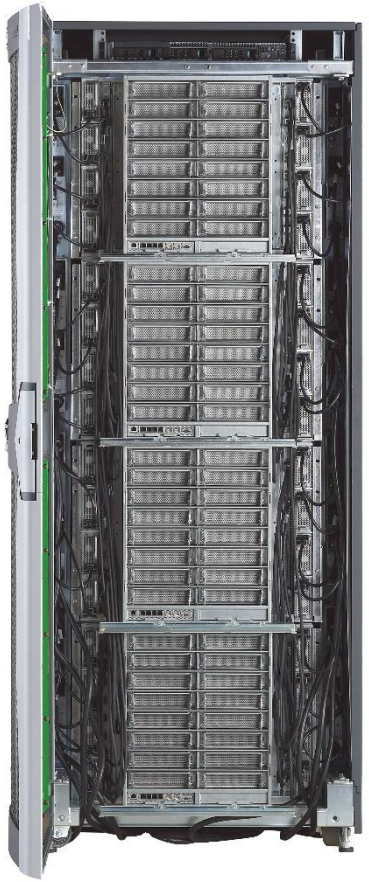

SGI ICE X Compute Node with GPU

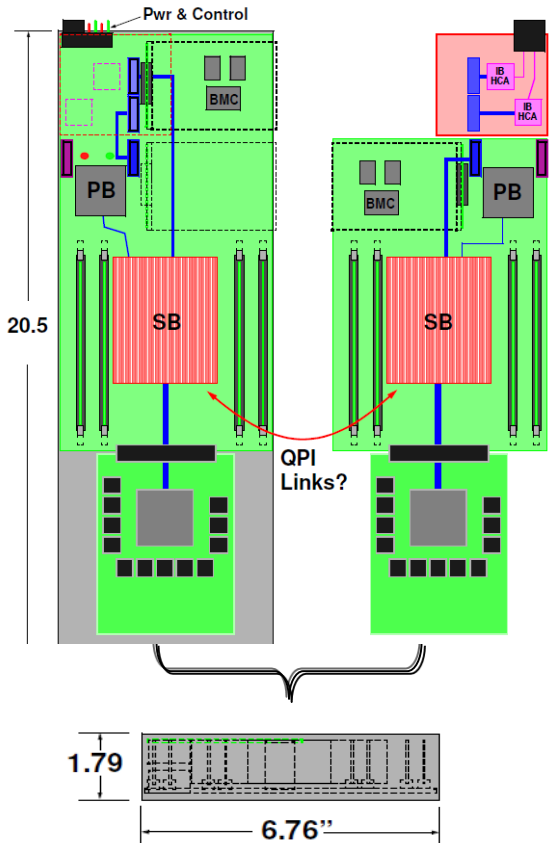

- Two Intel® Xeon® processor E5-2600 series
- **Two NVIDIA® TESLA® K20**
- FDR InfiniBand - single or dual plane
- Four DDR3 DIMMs per socket @ 1600 MT/s
- Up to one 2.5” SATA HDD/ SSD drives
- Liquid cold sinks

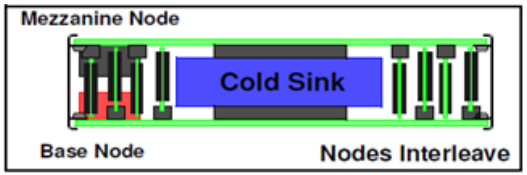

One dual socket node w/ **two** coprocessors in **one** blade slot!

FDR + 1:1 processor to coprocessor ratio = *Balanced Throughput*

19

SGI ICE X Compute Node with Phi

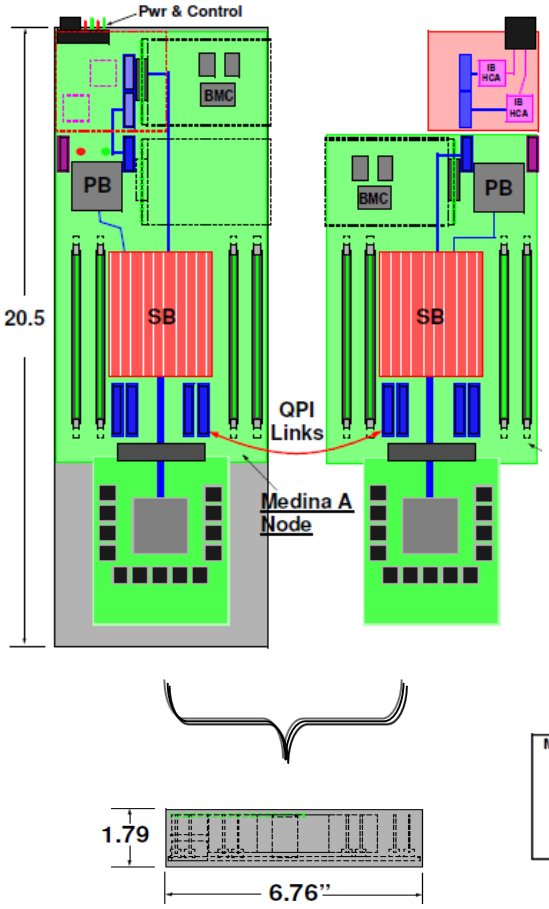

- Two Intel® Xeon® processor E5-2600 series
- **Two Intel® Xeon Phi™ 5120D Coprocessor**
- FDR InfiniBand - single or dual plane
- Four DDR3 DIMMs/ socket @ up to 1600 MT/s
  - Expect 1866 MT/s support w/1 DPC populated (w/ Ivy Bridge-EP)
- Up to one 2.5" SATA HDD/ SSD drive
- Liquid cold sinks

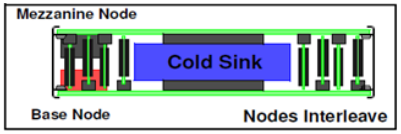

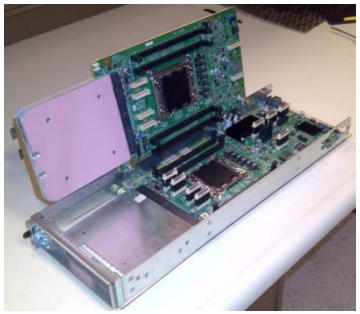

One dual socket node w/ **two** coprocessors in **one** blade slot!

FDR + 1:1 processor to coprocessor ratio = *Balanced Throughput*

SGI Cold Sink Technology – Twin Blades

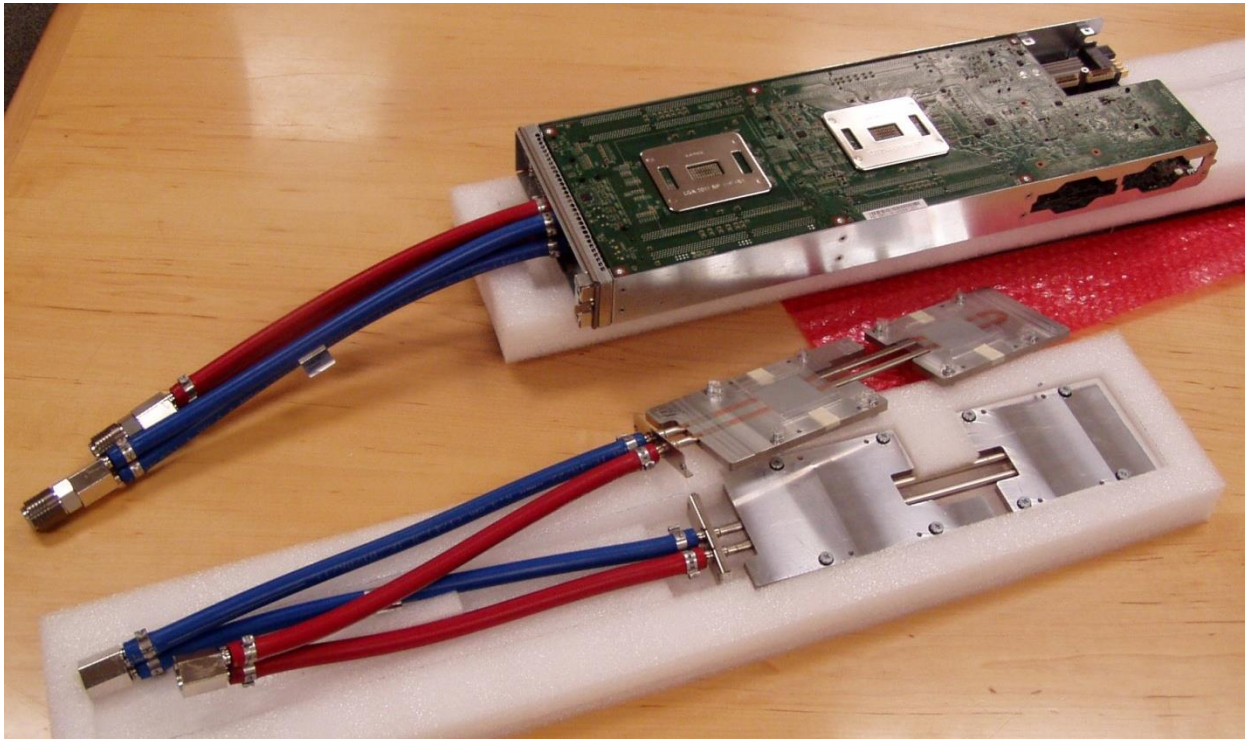

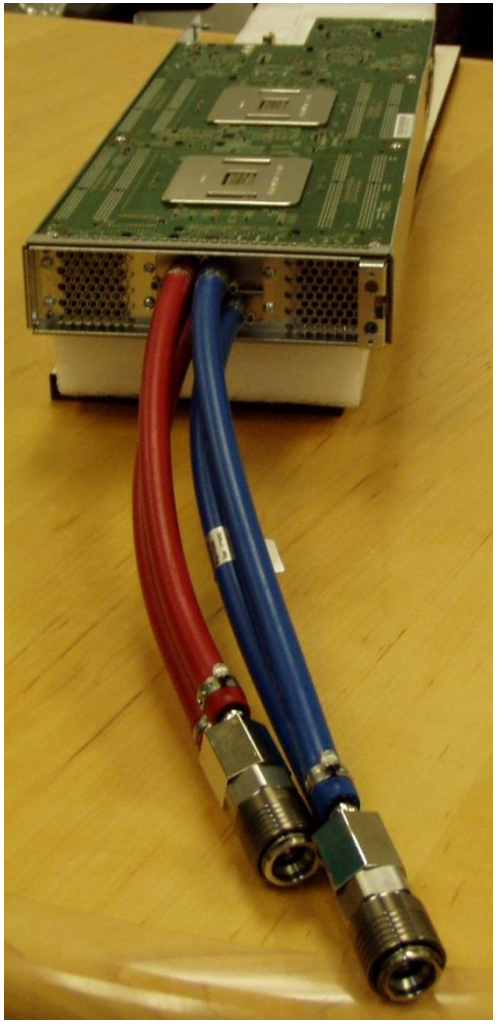

SGI Meets All Accelerator Needs!

- SGI has the longest history with accelerators
- Accelerators have gone mainstream, we are supporting them across our product line
- The only vendor to be able to deploy scale-up, scale-out and hybrid landscapes with accelerators everywhere.

22

SGI ICE: Sample ICE Customers

BAW - Bundesanstalt für Wasserbau	KU Leuven	Sikorsky
Bridgestone	Mazda Motor Corporation	Skoda Auto
Central Research Inst of Electric Pwr Ind	McLaren Motor Racing Ltd	TI-09
CSIR - Centre for Mathematical Modeling and Compute	Mercedes-Benz GP	Tokyo University – ISSP
Exeter University	MPO	Total
GENCI	NASA Ames	Toyota
HLRN Berlin	NASA Langley	U.S. Air Force - Arnold AFB
HLRN Hannover	NIIFI	U.S. Navy/NRL
ICHEC	NIMS	Universidade Catolica
ICR	NMCAC	Universite Paul Sabatier – CICT/Calmip
Idaho National Laboratory	NOAA (CSC)	University of Arizona
IFREMER	NTNU	University of Hyderabad
Imperial College	Onera	University of Oxford
IMSc	ORNL	University of São Paulo – USP/IAG
INMET	Pontificia Universidade Catolica – PUC-Rio	University of Rio de Janeiro – NACAD
JAMSTEC	Rosshydromet	University of Rio Grande do Sul – CESUP
Korean Air Force	Scientific Analysis Group	US Army TACOM
	Semiconductor Energy Labs	

Hungary 2011-2014

NIIFI supercomputers for non-profit customers:

- Cluster: SGI ICE
  Debrecen (18TFlops, 1536 core; 6TB ram)
- Cluster + GPU: HP
  Szeged (14TFlops, 2304 core; 5,6TB ram, 6xM2070)
- SMP/ccNUMA: SGI UV
  Pécs (10TFlops, 1152 core, 6TB ram)

http://www.niif.hu/szolgaltatasok/szuperszamitastechnika/altalanos_ismerteto

Hungary 2015

- **Cluster: 200+ Tflops**
  (200pc GPU or PHI)

- **Cluster: 30+ GPU or PHI**

- Cluster: SGI ICE
  Debrecen (18TFlops, 1536 core; 6TB ram)
- Cluster + GPU: HP
  Szeged (14TFlops, 2304 core; 5,6TB ram, 6xM2070)
- SMP/ccNUMA: SGI UV
  Pécs (10TFlops, 1152 core, 6TB ram)