Cyclone™
SGI Cloud Computing for HPC

Christian Tanasescu
Vice President Software Engineering

Agenda

- Rationale for Cyclone SGI offering
- Role in SGI business model
- Cyclone service and usage models
- Partnerships
- Future Directions

Limitations of Existing Cloud Services for HPC

- Most cloud systems built on virtualized instances only
- Limited scalability of available servers:
  - Scale-out clusters only
- Lack of high speed networks
  - Mostly GigE
- Lack of user control over node interconnect topology
  - MPI latency, core distribution, contiguous memory
- Lack of available HPC software stacks
- Lack of technical application specific configurations

Benefit of HPC Cloud Computing

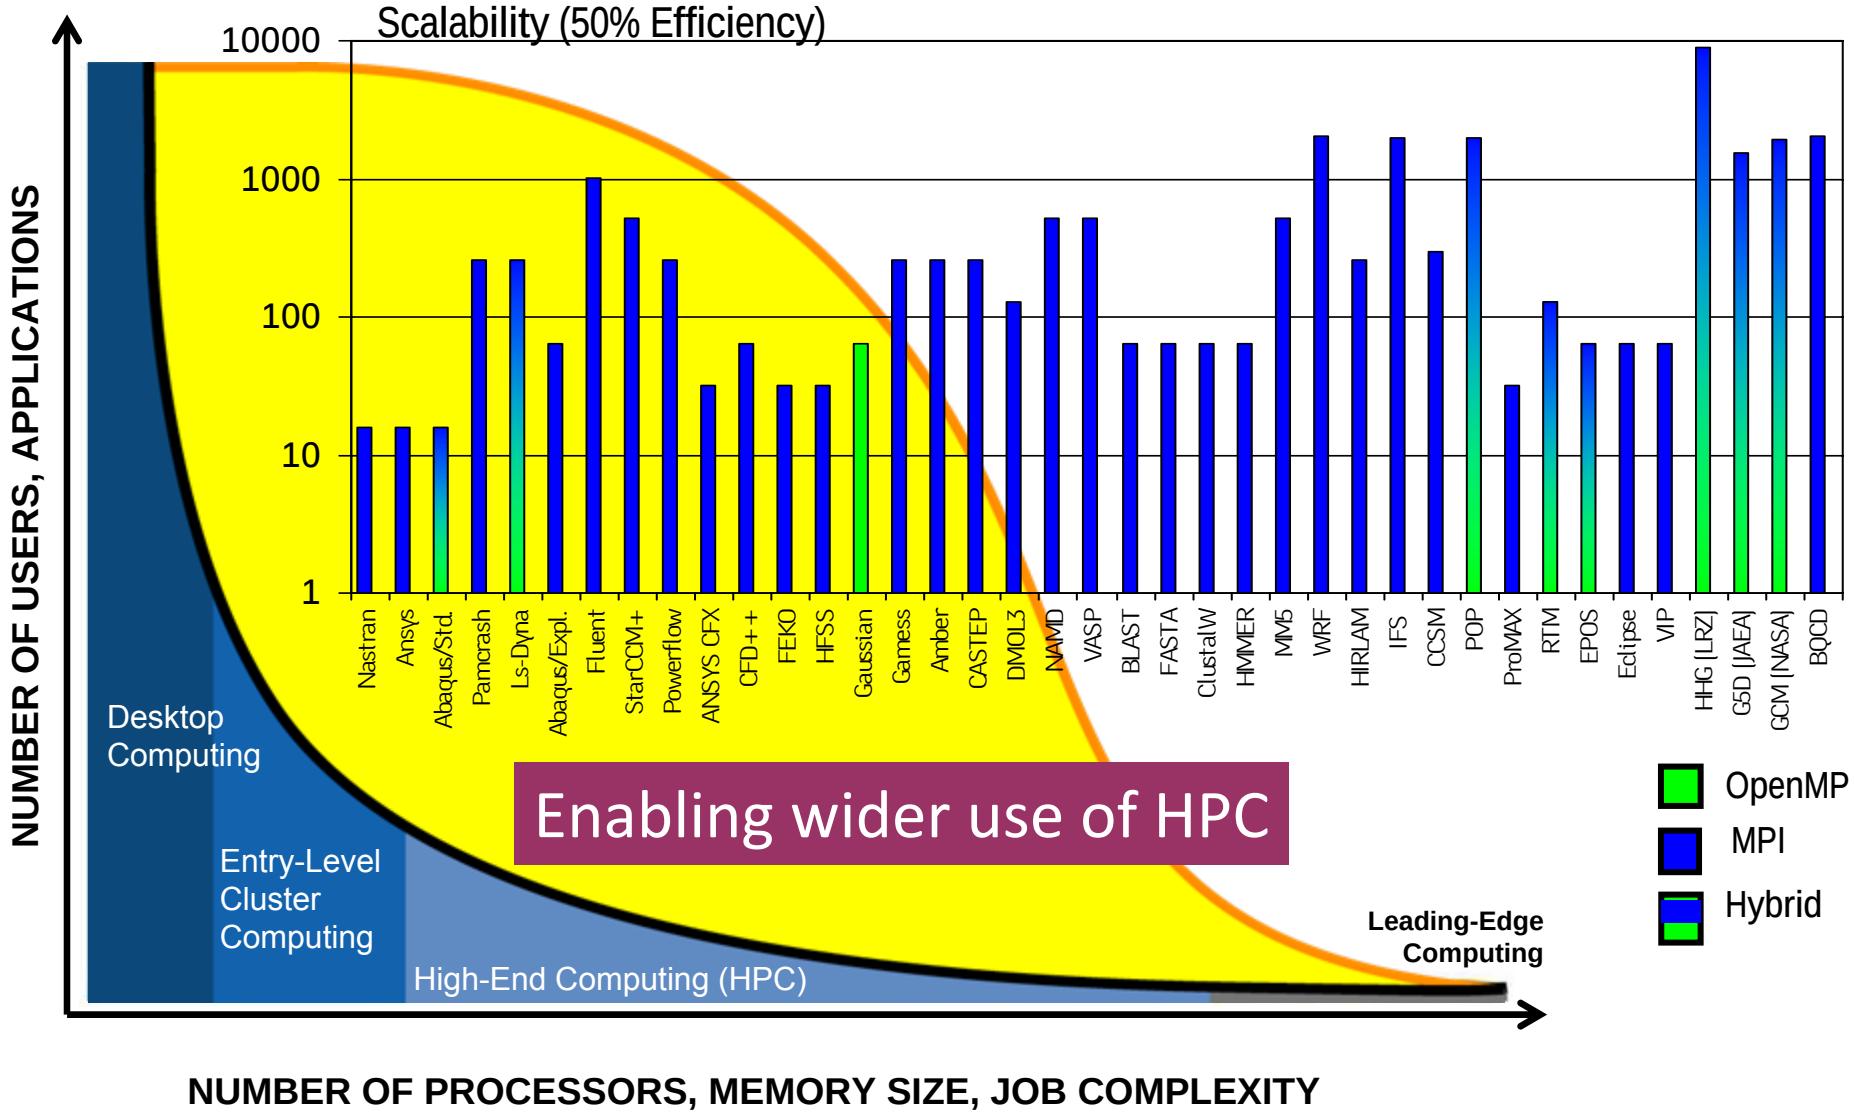

Source: Nimbis Services, 2010
Graphic adopted from OSC, Council on Competitiveness and the University of Southern California.

Completing the need

Personal& Workgroup	Traditional Data Center	Modular Data Center	Cloud
			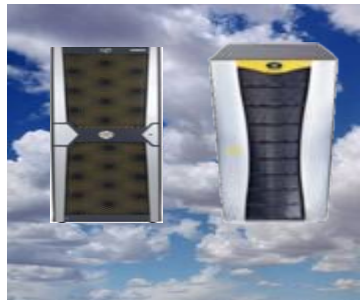
Access& Compute	Scale & Control	Modular & Mobile	On Demand

Move freely between the environments

SGI Cyclone- results on demand™

- Cloud computing dedicated to technical computing
- Cyclone Offers Flexible Choice of
  - Platform
    - Scale-up (Altix 4700 and Altix UV) and Scale-out (Altix ICE and Altix XE)
    - Hybrid systems with NVIDIA Tesla, ATI FireStream and Tilera accelerators
  - Operating System (SLES, RHEL)
  - Interconnect (NUMAlink, InfiniBand, GigE)
  - Physicalization and virtualization
- Application-Specific Cloud
  - Application-tuned software stacks
  - Open source and close source applications
- Differentiators
  - Significantly broader platform flexibility
  - SGI Performance Suite for application acceleration
  - Deep Application Engineering expertise

SGI Cyclone - Usage Models

- **Cloud Service for Customers**
  - SaaS and Iaas
  - Bridge for overflow capacity
  - Bridge for new system installation
  - Hub for new applications and architectures
- **Cloud for Software Development**
  - SGI Solutions Partner Program
- **Internal Cloud**
  - Benchmarking
  - Software development

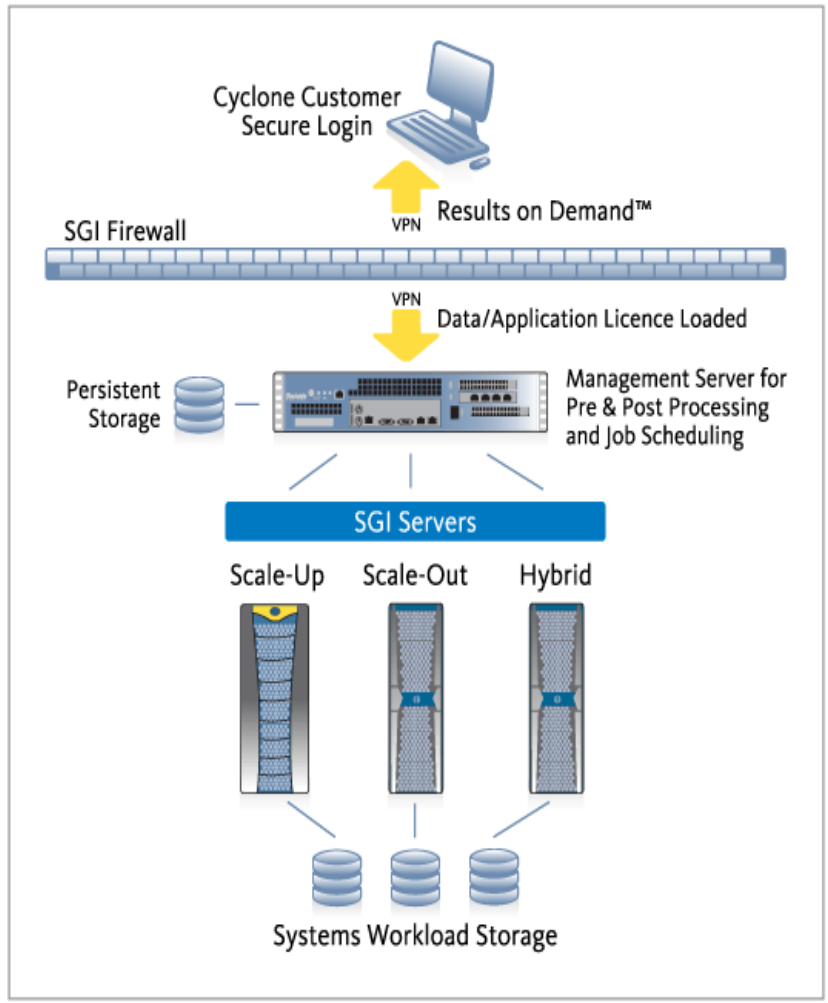

SGI Cyclone Domain and Application Focus

Domain	Use Cases	Applications
Computational Biology (BIO)	Genomics, proteomics, molecular modeling and drug discovery	BLAST, FASTA, HMMER, ClustalW
Computational Chemistry and Materials (CCM)	Nanotechnology, materials research (metal alloys, polymers, composites, ceramics and plastics)	Gaussian, Amber, Gamess, Namd, Gromacs, LAMMPS, DL_POLY
Computational Fluid Dynamics (CFD)	Automotive design, aerospace design, defense systems design and power generation design	StarCCM+, OpenFOAM, Acusolve, NUMECA
Computational Structural Mechanics (CSM)	Structural, heat transfer, fatigue and vibrational analysis for manufactured products	LS-DYNA
Computational Electromagnetis (CEM)	Design and analysis of antennas, antenna placement, EMC (shielding, coupling.), RF components and bioelectromagnetic analysis	FEKO
Ontologies (ONT)	Semantic Web, data mining	OntoStudio, SemanticMiner

Six Initial Domains: More to follow

19 Supported Applications: More to follow

SGI Cyclone System Architectures

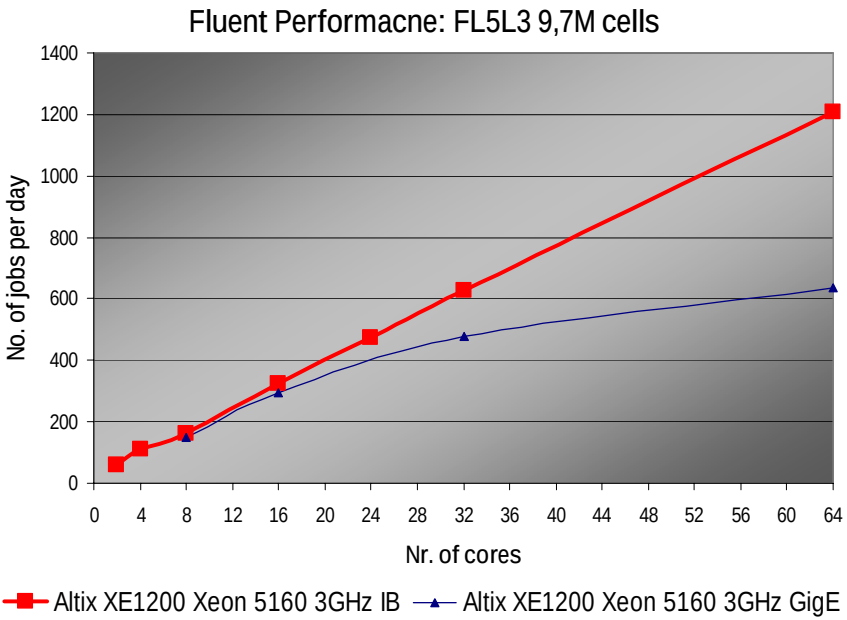

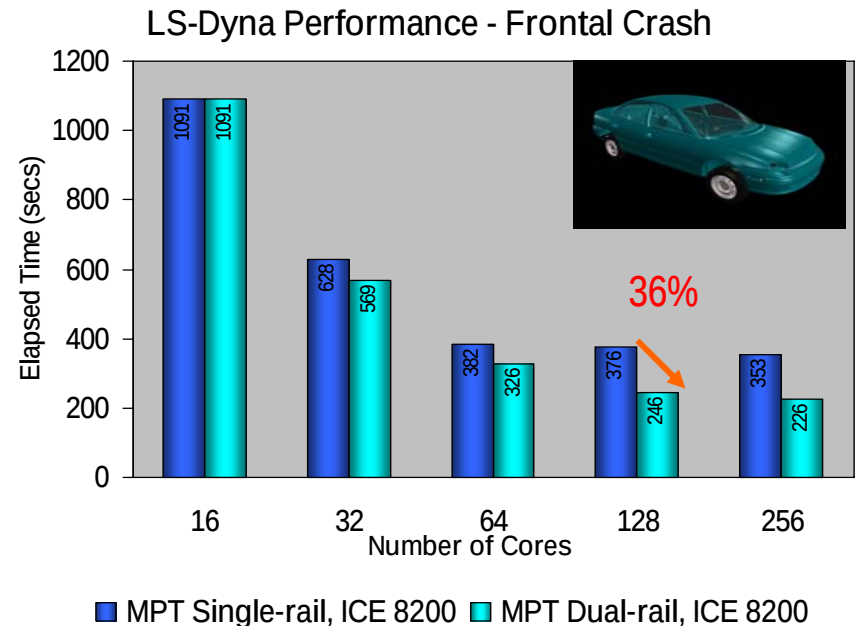

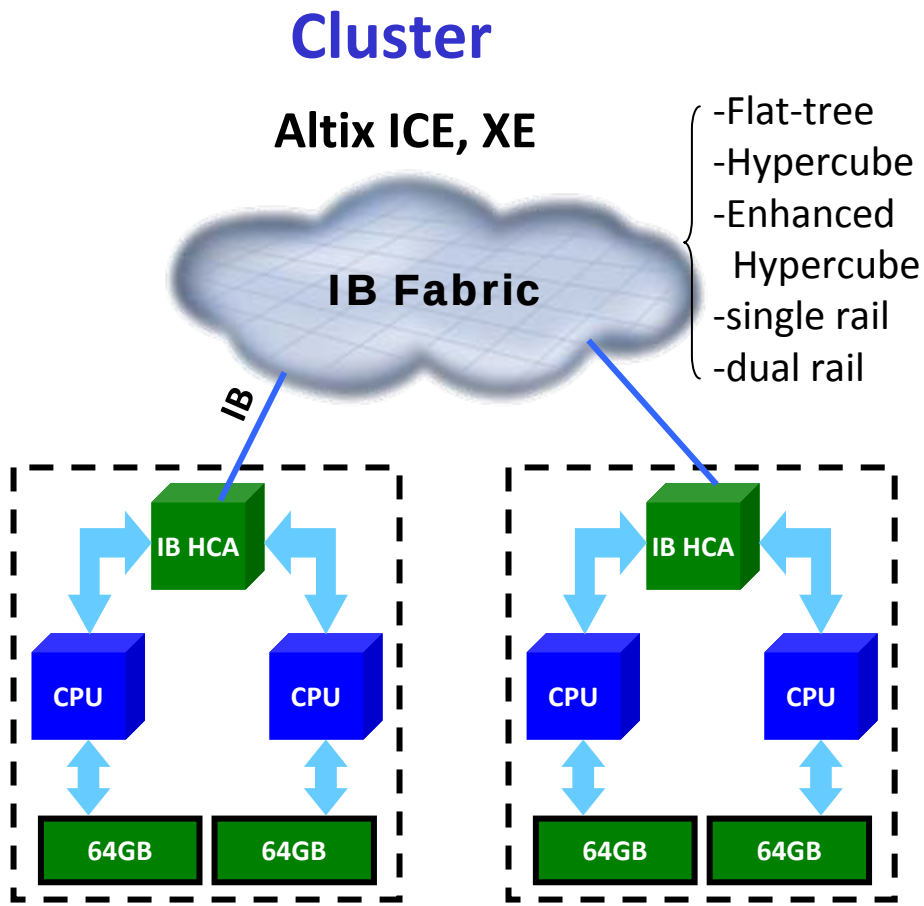

Platform interconnect choices determine the achieved application perforamance

SGI Cyclone System Architectures

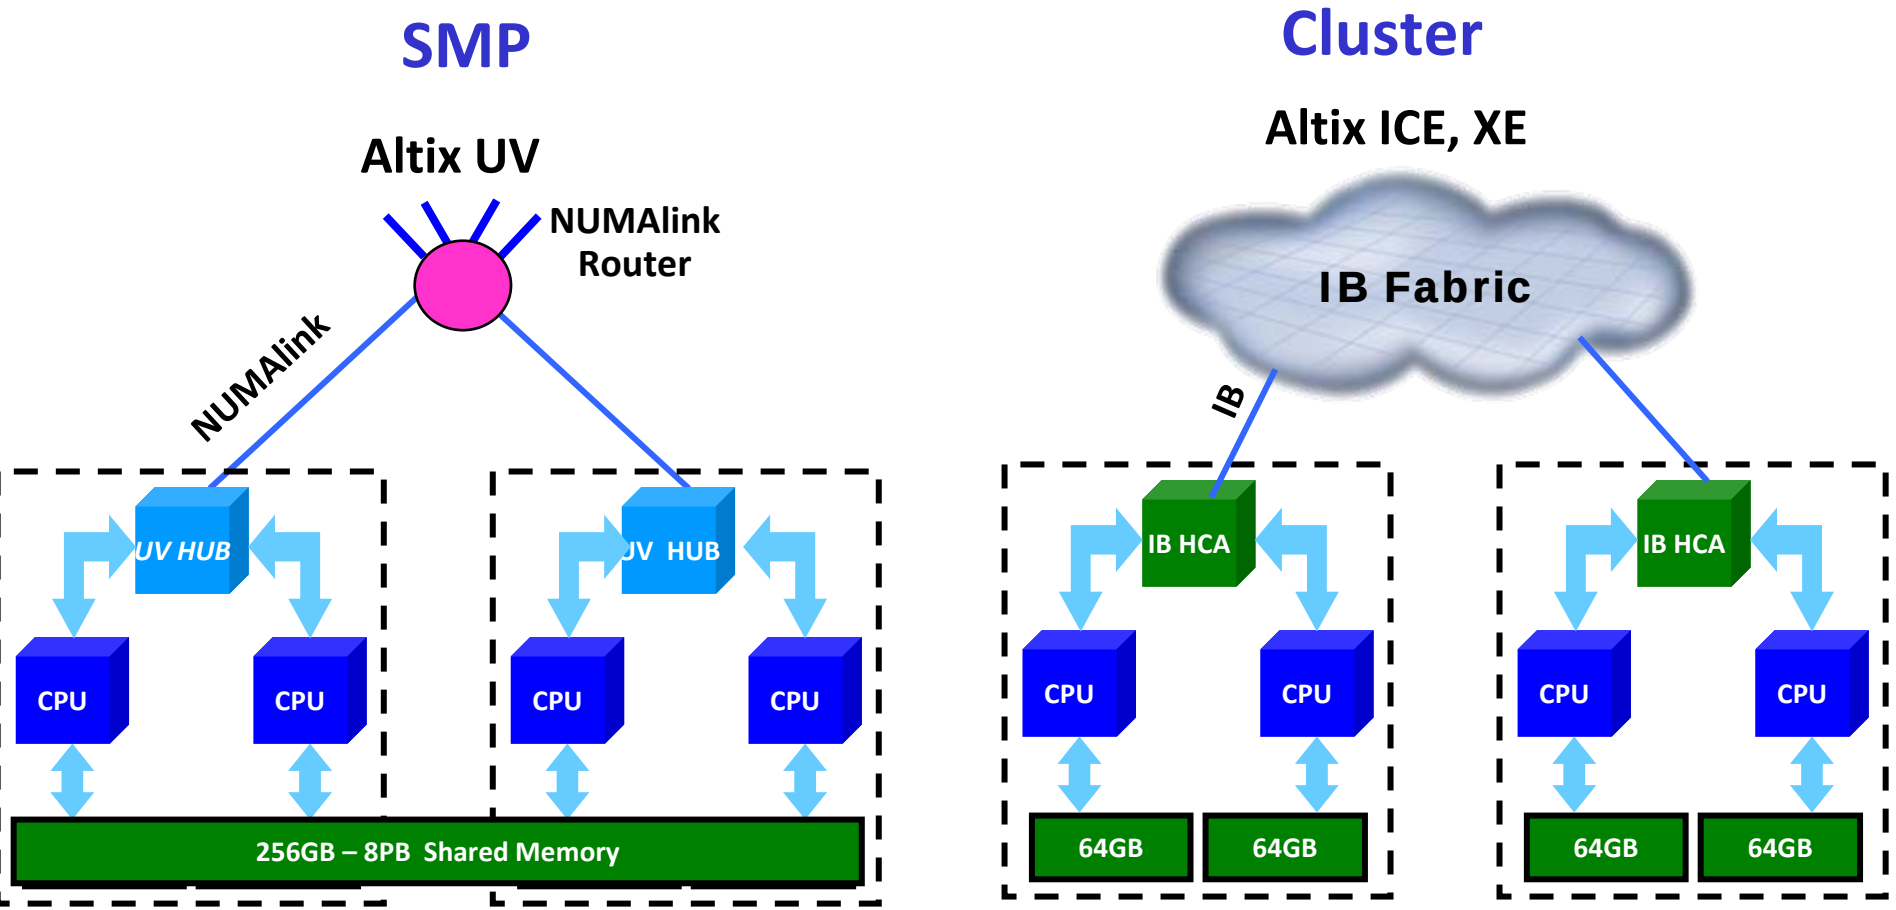

- NUMAlink® 5 is the glue of Altix® UV
- Global Shared Memory
  - Sockets: 2 to 256 in in one OS image system
  - Memory: up to 16TB in one OS image system
  - NUMAlink5 BW: 15 GB/s aggregate

- MPI off-load engine on Altix® UV
  - MPI reduction, barrier -> Synchronization 3x - 10x
  - MPI short messages -> local Latency 2x
  - MPI next neighbor communication -> local BW 3x
  - MPI Barrier Latency <1usec (4096 thread)

Performance Preview on Altix UV

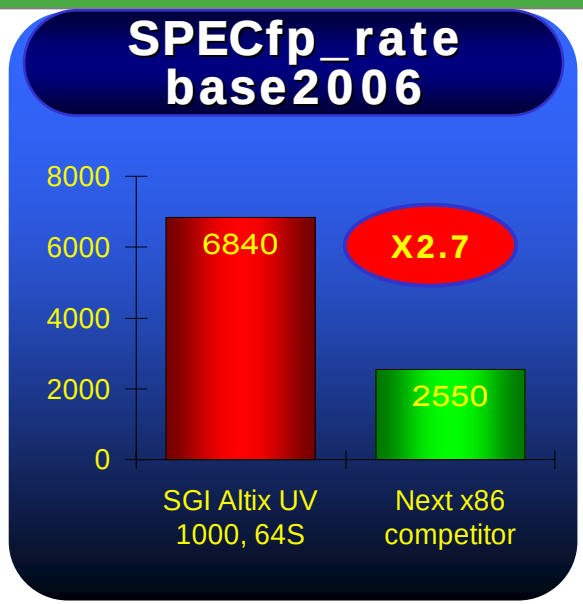

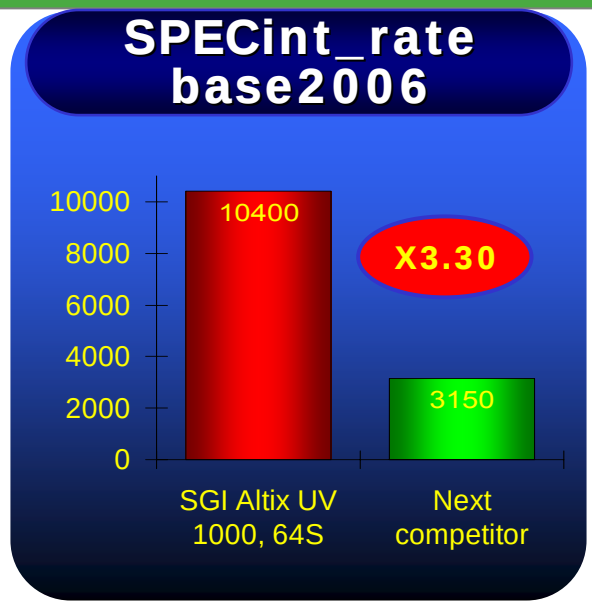

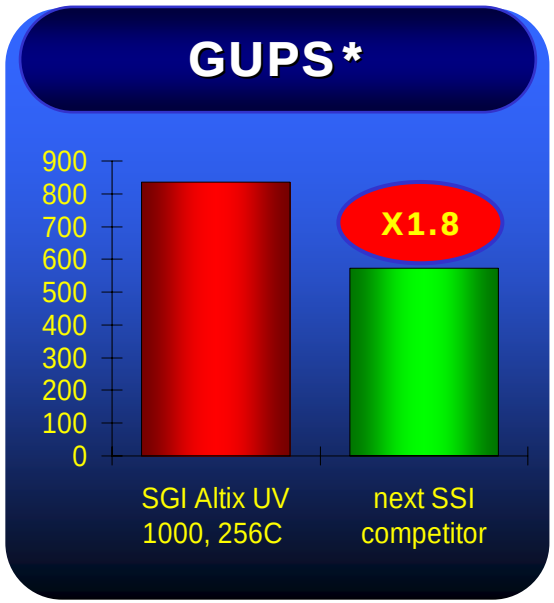

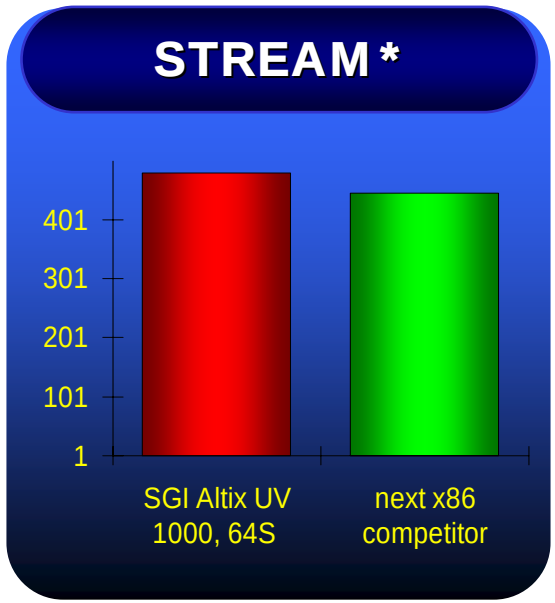

SGI Altix UV is the industry's leading supercomputer in relevant performance metrics

sgi

Rationale for GPU
Analytics of Top500 list, Nov 2009

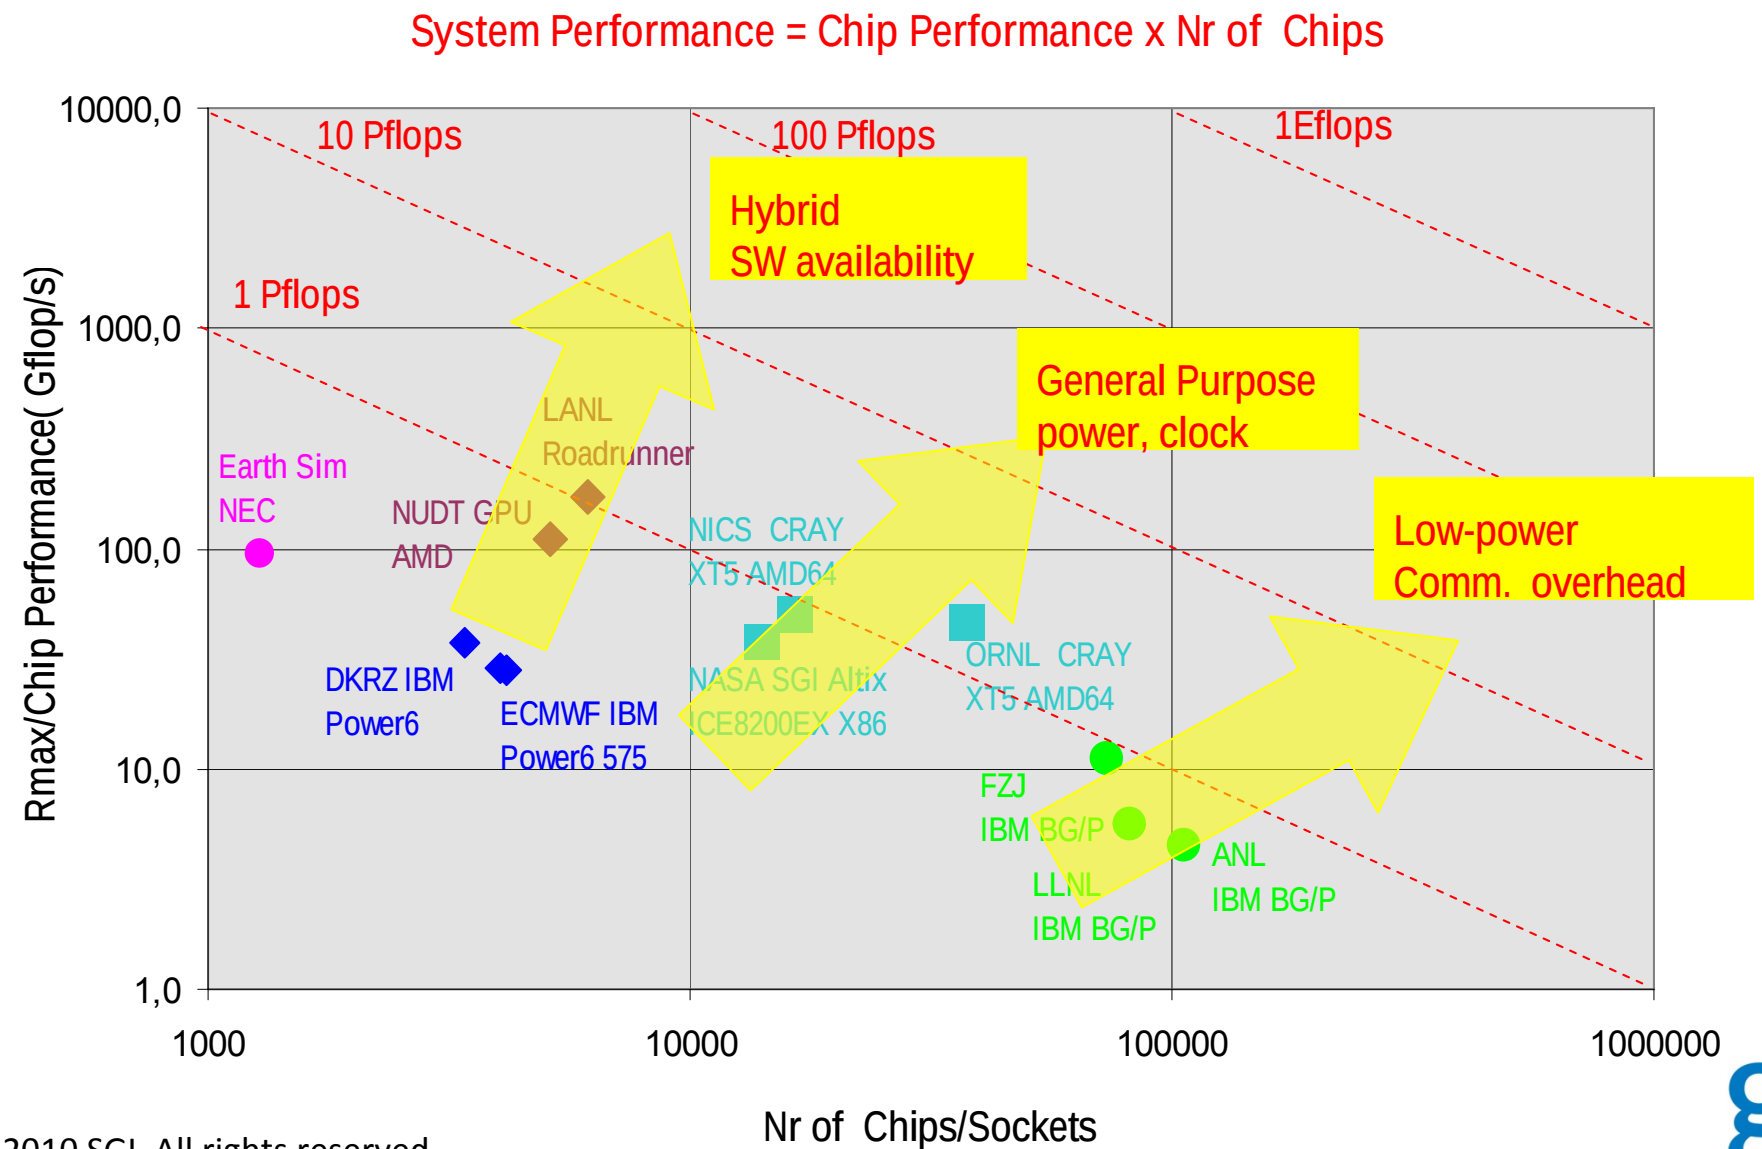

SGI Cyclone Software Stack

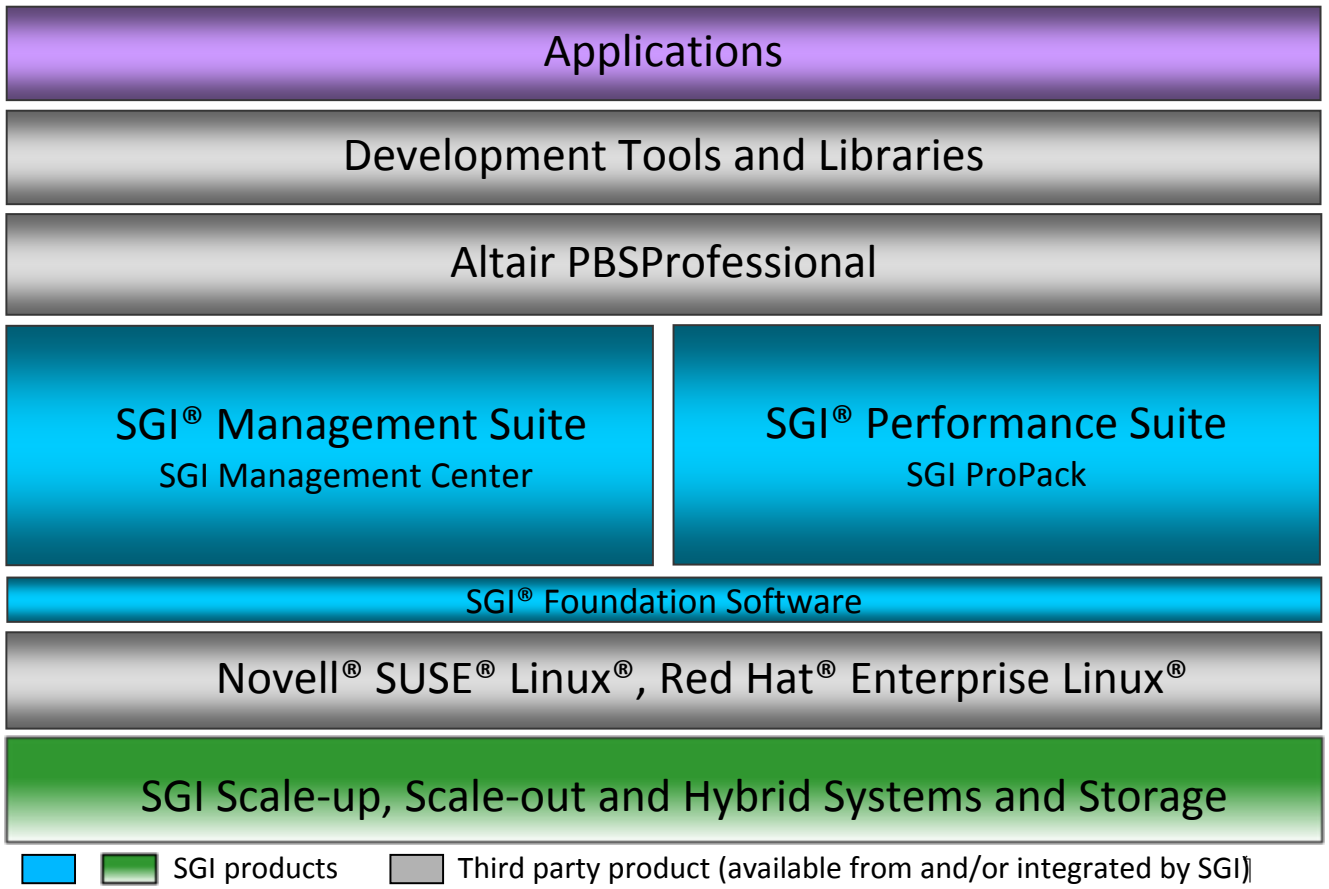

SGI Performance Suite – Application Accelerator

FFIO – IO Accelerator
linkless IO library that enhances performance for jobs that have much larger IO footprint sizes than free memory sizes on the systems.

SGISolve
-Scalable SMP parallel in-core and out-of-core sparse solvers
- Parallel iterative solvers and pre-conditioners library

MPT - MPI offload engine for SMP and Clusters

PerfBoost– Accelerate applications certified HP-MPI, Intel MPI or OpenMPI No recompile or re-linking needed

XPMEM, XPNET – fast cross-partition data transfer

MPInside – MPI performance modelling and prrojection tool

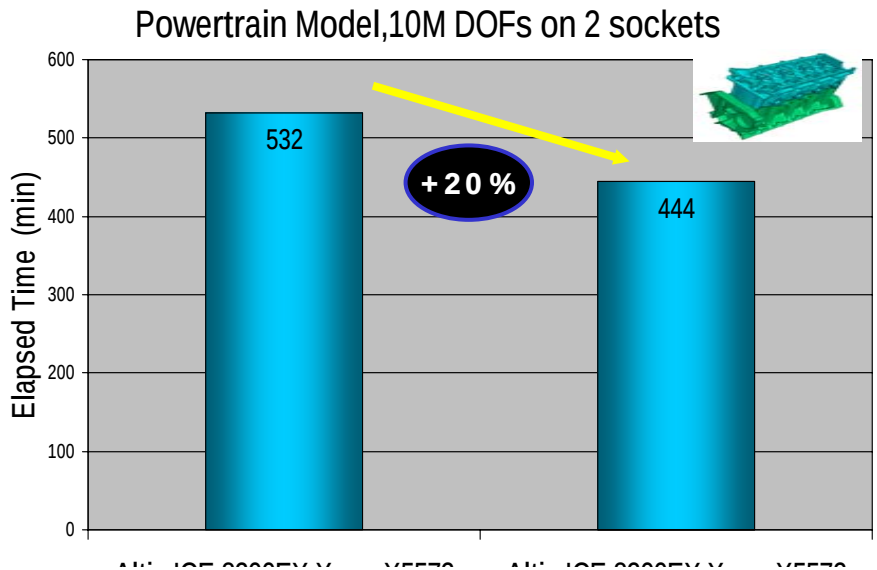

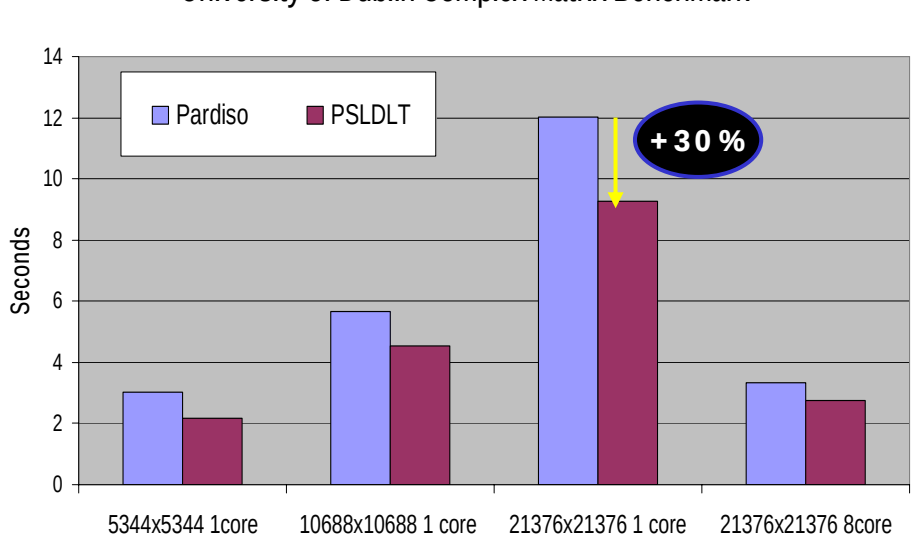

Partnership - Nimbis Cloud Portal for SGI Cyclone

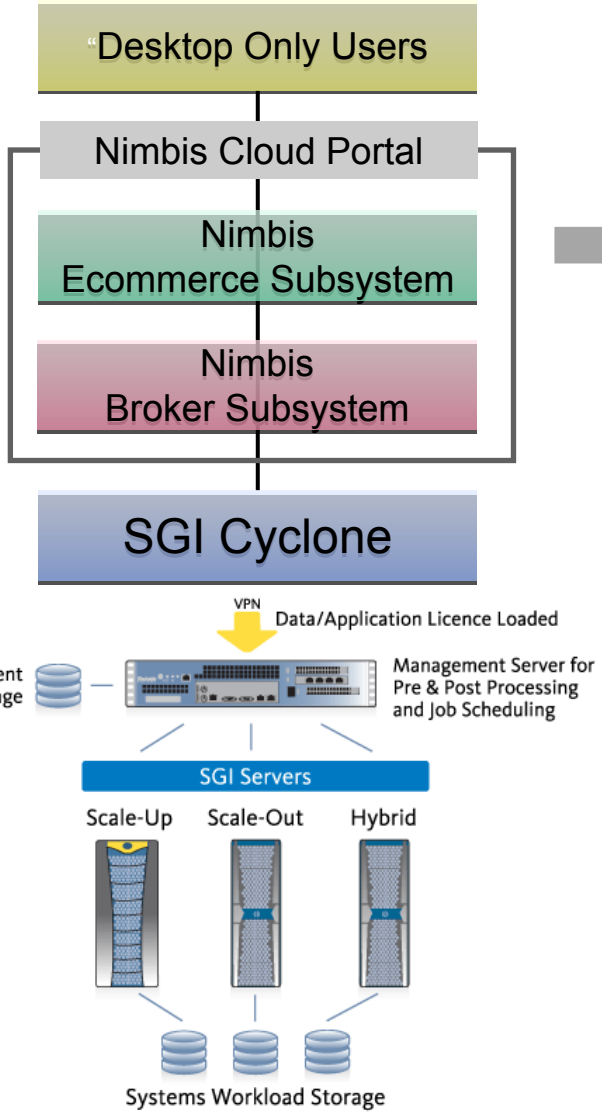

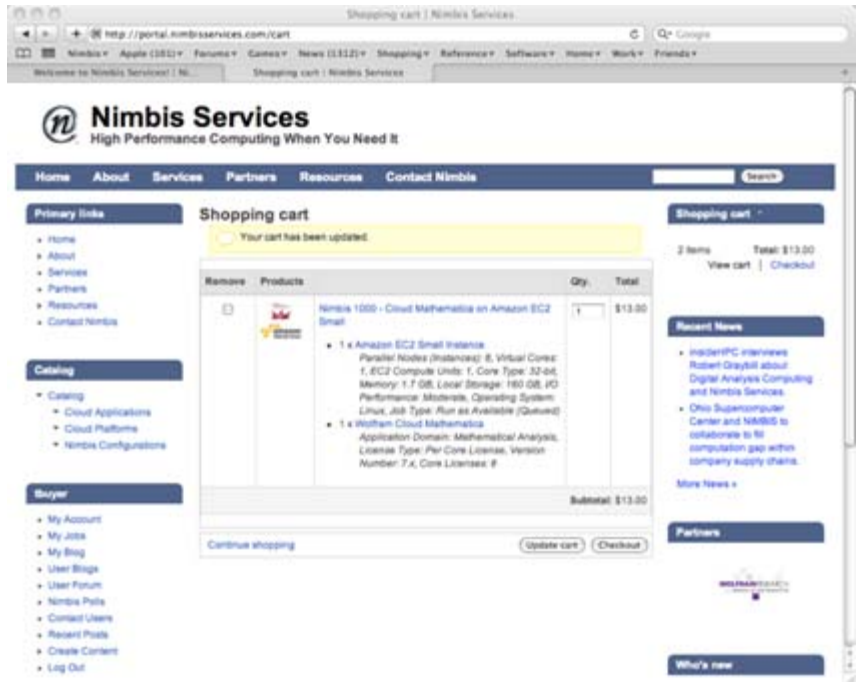

Nimbis Services

Cyclone Technical Application Library
Application-tuned HPC environments

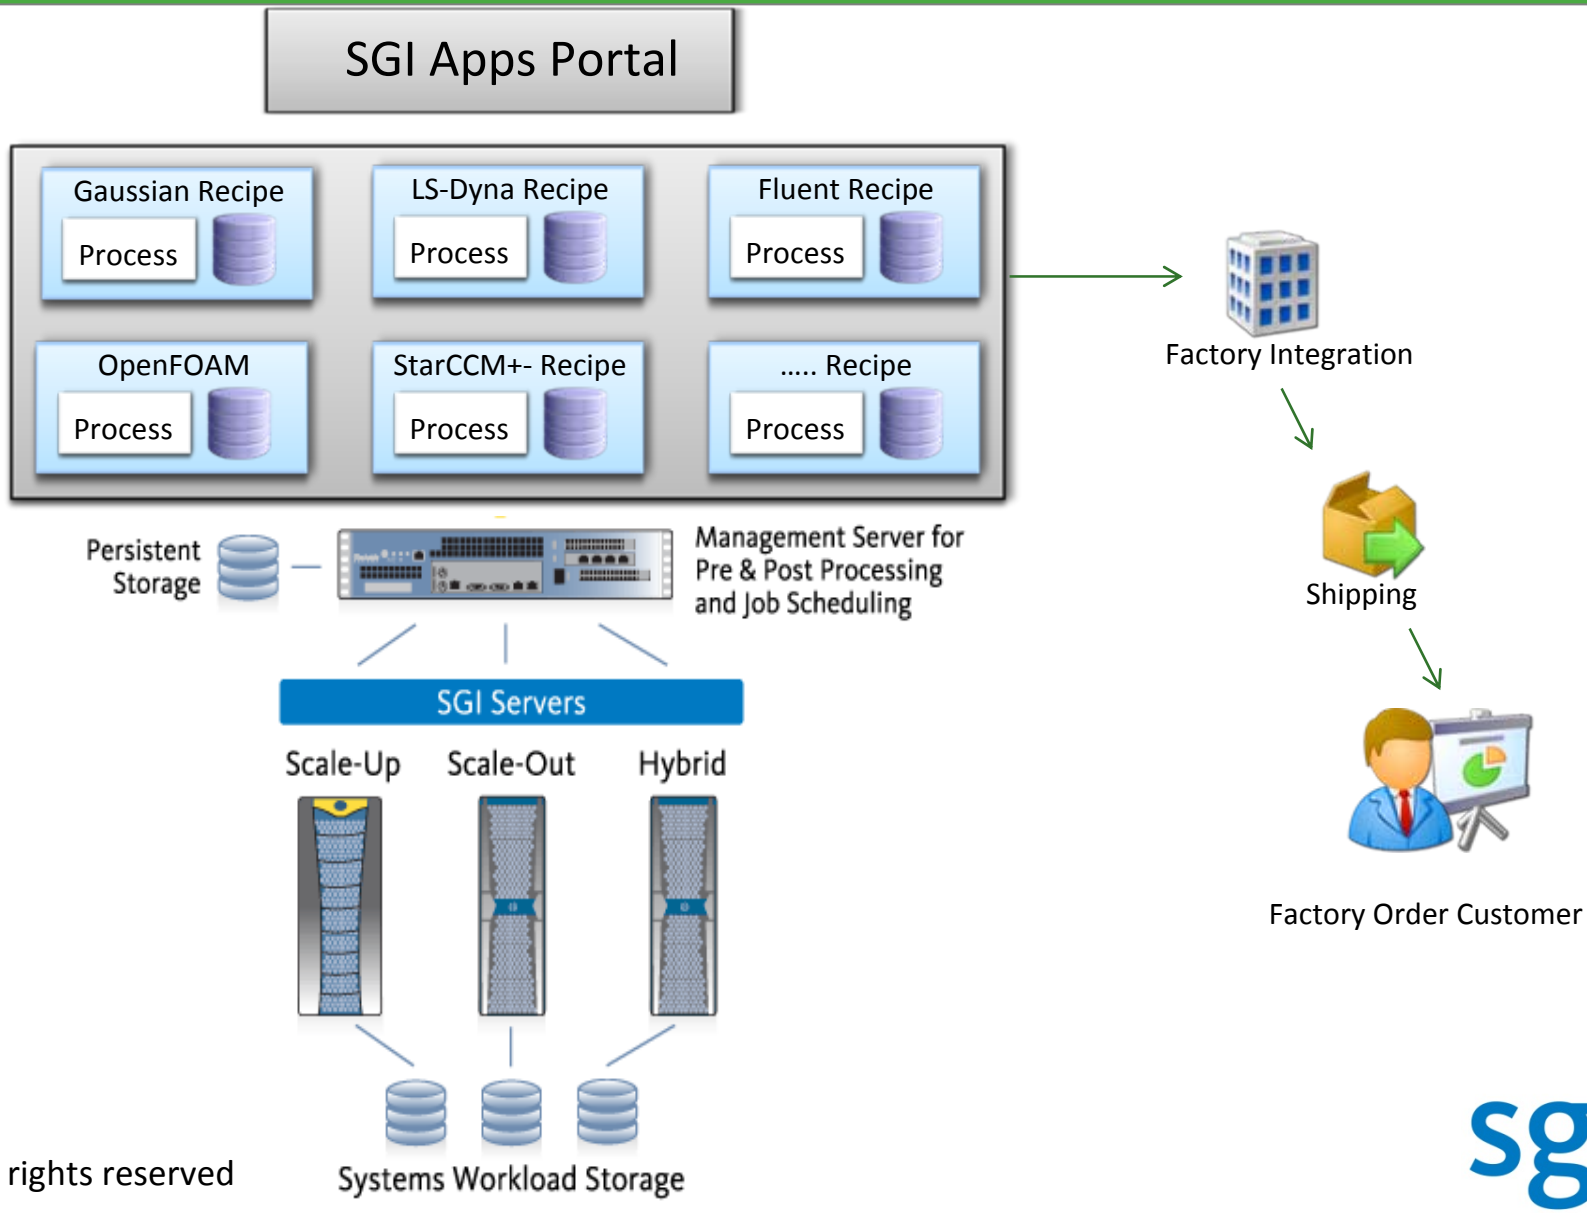

Conclusions

- Cyclone completes the need
- Specifically dedicated to technical computing
- Making HPC pervasive
- Fast access to new technologies