# Scalable x86 SMP Server

# **FUSION1200®**







# Challenges

## Scaling compute-power is either Complex (scale-out / clusters) or Expensive (scale-up / SMP)

## Scale-out - Clusters

- Requires advanced IT skills / know-how (high OPEX):
  - <u>Complex</u> operating environment: multiple OS, Cluster file-system, distributed software installation, cabling...
- Suitable for large deployments (>64-processors)
  - OEM operates on <u>thin margins</u> (<20%)</li>





## Scale-up - SMP (shared memory systems)

- High end-user cost Per processor price grows with system size
  - High OEM <u>R&D expenditure</u> to create x86 SMP systems (>4 processors)
  - Significant margin expansion for OEMs
- Product offering is mostly proprietary (non-x86) in nature
  - Architecture and processors (Itanium, SPARC, POWER, PA-RISC...)
  - Limited application portfolio





## ScaleMP's Approach: A3xs

ScaleMP's versatile SMP (vSMP) architecture:

- Aggregates standard components to larger system
- Scales application processing, memory and I/O resources
- Simplifies system management by reducing number of systems
- Saves through lower TCO



- x86-based shared memory systems
- Off-the-shelf components
  - Reduce R&D cost
  - Best price-performance
- Higher Margin offering





# vSMP Foundation – How It Works?



High-end x86 system, based on standard x86 components





## vSMP Foundation Configuration Versatility



Flexible high-end x86 system, meeting end-user requirements





# vSMP Foundation – Behind The Scenes

## **One System**

- Software interception engine creates a uniform execution environment
- vSMP Foundation creates the relevant BIOS environment to present the OS (and the SW stack above it) as single coherent system

#### **Coherent Memory**

- vSMP Foundation maintains cache coherency between boards
- Multiple concurrent memory coherency mechanisms, on a per-block basis, based on real-time memory activity access pattern
- Leverage board local-memory for caching

#### Shared I/O

- vSMP exposes all available I/O resources to the OS in a unified PCI hierarchy
- No need for cluster file systems





# **Highest x86 SMP Memory Bandwidth**



**HW Characteristics:** vSMPowered System IBM x460: SUN X4600M2:

#### Description

32 x Intel Xeon QC 5345 – 2.33GHz, 2x4MB L2, 900/960GB 8 x Intel Xeon DC 7040 – 3.0GHz, 2x2MB L2, 48GB 8 x AMD Opteron DC 8218 – 2.6GHz, 2x1MB L2, 32GB ScaleMP ScaleMP ScaleMP



# **Stellar Application Performance(\*)**





(\*)More application performance information available



## FUSION1200®



- TODAY'S WORLD'S LARGEST INTEL<sup>®</sup> XEON <sup>®</sup> SMP SERVER
- OVER 70 SYSTEMS DEPLOYED
   WORLDWIDE
- ALL SYSTEMS ARE RUNNING
   IN PRODUCTION

## **Road To Success**





## FUSION1200®

## **128 CORES / 1TB SMP AT MANAGED CLUSTER PRICES!**





# **FUSION1200® – Competitive Overview**

System Vendor	VXTECH	IBM	SUN	White-Box		
System Model	FUSION-1200	x3950 M2	x4600 M2	AMD		
Processor						
Vendor	Intel	Intel	AMD	AMD		
Model	Xeon 52xx / 54xx	Xeon 72xx / 73xx	Opteron 8xxx	Opteron 8xxx		
Micro-Architecture	Intel Core	Intel NetBurst	AMD K8	AMD K8		
Dual-Core availability	Yes	Yes	Yes	Yes		
Quad-Core availability	Yes	Yes	Yes	Yes		
Chassis						
Max. Processors (sockets) / Chassis	12	4	8	8		
Max. Memory / Chassis (GB)	192 or 384	128	256	256		
Max. Internal Drives / Chassis	6	6	4	8		
Integral GigE Ports / Chassis	7	2	4	3		
System						
Max. Chassis	4	8	Not-Expandable	Not-Expandable		
Max. Processors (sockets) / System	32	16	8	8		
Max. Memory (GB) / System	512 or 1TB	512	256	256		
Max. Internal Drives / System	16	48	4	8		
Integral GigE Ports / System	20	16	4	3		
Others						
Internal Drives Type	SATA	SAS	SAS	SATA / SAS		
Integral IO Expansion / Chassis	Yes (eSATA)	No	No	No		
Available PCIx/e Slots / Chassis	1	6	8	5		
Redundant Power Supply	No	No	Yes	Yes		



## **Why Customers Buy?**

## **Application requirements**

- Applications that use large memory footprint (even with one processor)
- Applications that need multiple processors and shared memory

#### **FUSION system delivers:**

- Best price/performance
- For customer looking for x86, the only solution that delivers large memory and high-socket count

## **Operational requirements**

- Use same computing resource for multiple application classes ("mixed environments"): singleprocessor, multi-processor, sharedmemory
- Simplify the complexities involved with running clusters

## **FUSION system delivers:**

- The most flexible platform top performance for all application classes
- Lowest TCO!



# Finding opportunities for the FUSION1200®



#### **End user requirements**

#### Fusion1200 fit

**Key CAE messages** 

- **1)** Need more compute than current workstations
- 2) Run mix of Structural, CFD and Crash applications
- 3) Don't want clusters
- 4) Need to run more accurate models, faster

- 1) Keep workstation for graphics add shared server
- 2) Large memory for structural, many processors for CFD and crash
- 3) System administration equivalent to a workstation

1) Run multiple jobs in parallel (i.e. 2 jobs on 4 CPUs

3) Great system for FSI (large memory for structural,

high memory bandwidth and compute for fluid)

2) Run small jobs during day, large jobs at night

Six key differentiators (mid-range x86 SMPs)

4) Grow as needed: up to 32 CPUs -512GB RAM

each) or one large job on all 8 CPUs

4) Flexibility, flexibility, flexibility

5) SGI has applications expertise

Key Fusion1200 messages

3) Highest memory bandwidth

1) Highest socket count

5) Competitively priced

6) Ease of Management

2) Largest memory

#### Where to find opportunities

#### Companies that:

- 1) have several engineers, little IT staff or expertise
- 2) can share a common resource
- 3) need more RAM than the 64GB or 128GB limitations
- 4) run structural, CFD and Crash applications
- 5) want both x86 and SMP servers
- 6) need to generate very large meshes

#### Applications Tested in ScaleMP labs

- 1) ANSYS Mechanical
- 2) ABAQUS/Explicit & Standard
- 3) LSTC-LS-DYNA
- 4) FLUENT
- 5) ASNYS CFX
- 6) CD-adapco Star-CD
- 7) AVF FIRE

Many Benchmarks

Available



# **Life Sciences**

End user requirements	Fusion 1200 fit	
<ol> <li>Mix of Legacy OpenMP, MPI, and in-house applications (most of them x86)</li> <li>Run multiple applications on same system</li> <li>Support multiple users</li> <li>Performance is important, but not the only factor</li> <li>Need to publish results fast (no time for parallelizing code using message passing)</li> </ol>	<ol> <li>High number of processors well suited for many computational chemistry and bio-informatics codes</li> <li>High Memory bandwidth needed by several computational chemistry codes</li> <li>Ability to run hundreds of x86 life sciences applications (not all run on Itanium)</li> <li>SMP system, easier to install and administer</li> <li>Optimal Perf. Guidelines available for many apps.</li> </ol>	
Where to find f1200 opportunities	Key Life Sciences messages	
<ol> <li>Centers considering IBM 3950 and Sun X4600</li> <li>Sites that have experience with SMP systems, easier sell</li> <li>Little/no IT staff. Can't manage complex systems</li> <li>University research sponsored by grants (usually in the \$200K with 50 to 100K for HW)</li> <li>Bio-techs (either will buy white boxes clusters, or value the SMP)</li> <li>Larger Pharmaceuticals. Go to research department. Be aware of corporate IT preferred vendors.</li> </ol>	<ol> <li>Run ANY 86 application</li> <li>Run multiple jobs and one large jobs</li> <li>Run small jobs during day, large jobs at night</li> <li>Larger shared resource is more flexible</li> <li>SGI has strong industry expertise</li> <li>Great performance against IBM and Sun systems</li> <li>Performance information available on key applications</li> <li>Scalability, start with a chassis and grow over time</li> </ol>	
Applications Tested in ScaleMP lab	Key Fusion1200 messages	
Gaussian – Amber – GROMACS - Schrödinger Jaguar – Schrödinger Glide – NAMD- DOCK – GAMESS – GOLD – mpiBLAST - BLAST – MOLPRO - OpenEye FRED - OpenEye OMEGA - SCM ADF - HMMER Many Benchmarks	Six key differentiators (mid-range x86 SMPs) 1) Highest socket count 2) Largest memory 3) Highest memory bandwidth 4) Grow as needed: up to 32 CPUs -512GB RAM 5) Competitively priced 6) Ease of Management	
Available		



# **Numerical Simulations**

Fusion1200 fit		
<ol> <li>High number of processors and processor cores to run many MATLAB simulations in parallel</li> <li>High Memory capacity for very complex models, larger than any other x86 system</li> <li>Ability to support hundreds of users</li> <li>SMP system, easier to install and administer</li> </ol>		
Key MATLAB messages		
<ol> <li>Run 100's of jobs in parallel, one core per user/process</li> <li>Run jobs that require very large memory. A single process can user the entire memory of the system</li> <li>Run one job in parallel on entire system (requires StarP or MATLAB DCE)</li> <li>Run small jobs during day, large jobs at night</li> <li>Larger shared resource is more flexible</li> <li>SGI has strong industry expertise</li> <li>Scalability, start with a chassis and grow over time</li> </ol>		
Key Fusion1200 messages		
Six key differentiators (mid-range x86 SMPs) 1) Highest socket count 2) Largest memory 3) Highest memory bandwidth 4) Grow as needed: up to 32 CPUs -512GB RAM 5) Competitively priced 6) Ease of Management		



## **Targets For FUSION1200® Systems**

#### TARGET ENVIRONMENTS

- Applications that use **large memory** footprint
- Applications that need multiple processors and shared memory
- User wanting to simplify the complexities involved with running clusters

#### MANUFACTURING

Computational Structural Mechanics (CSM) • ABAQUS/Explicit - ABAQUS/Standard - ANSYS Mechanical - LSTC LS-DYNA - ALTAIR Radioss

Computational Fluid Dynamics (CFD) • FLUENT - ANSYS CFX - CD-Adapco STAR-CD - AVL FIRE

Others

inTrace OpenRT

#### LIFE SCIENCES

 AMBER - Schrödinger Jaguar - Schrödinger Glide -NAMD - DOCK - GAMESS- GOLD - mpiBLAST -GROMAX - MOLPRO - OpenEye FRED - OpenEye -OMEGA - SCM ADF - HMMER

#### ENERGY

 Schlumberger ECLIPSE - Paradigm Geophysical GeoDepth

#### **HORIZONTAL & SYNTHETIC BENCHMARKS**

- The MathWorks MATLAB SPEC CPU2000
- STREAM (OMP)



# **Selling Against The Competition**

Vendor	Products	Presence in Industry	How to Compete		
IBM	X3950/M2 P-Series	<ul> <li>Mostly found in life sciences</li> </ul>	<ul> <li>Poor application performance on x3950 due to limited memory bandwidth (1066MHz FSB)</li> </ul>		
		<ul> <li>No presence in manufacturing</li> </ul>	<ul> <li>FUSION is much stronger (1333MHz x 6 FSB)</li> <li>Be aware of preferred vendors in life sciences</li> <li>FUSION beats P-Series in customer benchmarks</li> </ul>		
Sun	X4600	<ul> <li>Strong in Higher Education and Research</li> <li>Little presence in manufacturing</li> </ul>	<ul> <li>f1200 better performance, more scalability</li> <li>X4600 AMD systems do not scale past 4 processors and lack quad core offering !</li> </ul>		
НР	Does not have an offering in 8+ CPUs Shared Memory Processor Systems				
DELL	Does not have an offering in 8+ CPUs Shared Memory Processor Systems				



# **Handling Objections**

## UNPROVEN TECHNOLOGY

- 70+ systems deployed- sold systems in production
- Blue-Chip Repeat customer purchases
- Guaranteed better performance at end of life (vSMP Foundation firmware updates, with performance improvements)

## • **EXPENSIVE**

- Competitively priced vs. other x86 SMPs
- Small premium over enterprise class managed clusters

## NEEDS TUNING AND APPLICATION SET UP

- All applications run unmodified
- Runs standard Linux distributions
- As in any architecture, guidelines apply for best performance

## TAKES LONG TIME TO PRICE

Online Configurator available for rapid configuration pricing



# **Engagement Model**



