Building HPTC Solutions from a Solution Architect's View

Dave Mullally, Ph.D.
ISV Solutions Engineering
Curtis Bennett
HPTC Americas SA
Joel Jaffe
HPTC Americas Solutions Design

Session #1353

HP WORLD 2003
Solutions and Technology Conference & Expo

Today's Agenda

▪ Introduction	Joel Jaffe
▪ Life Sciences Perspective	Dave Mullally
– Overall Market	
– Computational Chemistry	
– Bioinformatics	
– Clinical Trials	
▪ CAE Perspective	Curt Bennett
▪ Summary and Q&A	Team

The Logical Clustered System

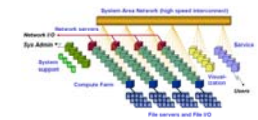

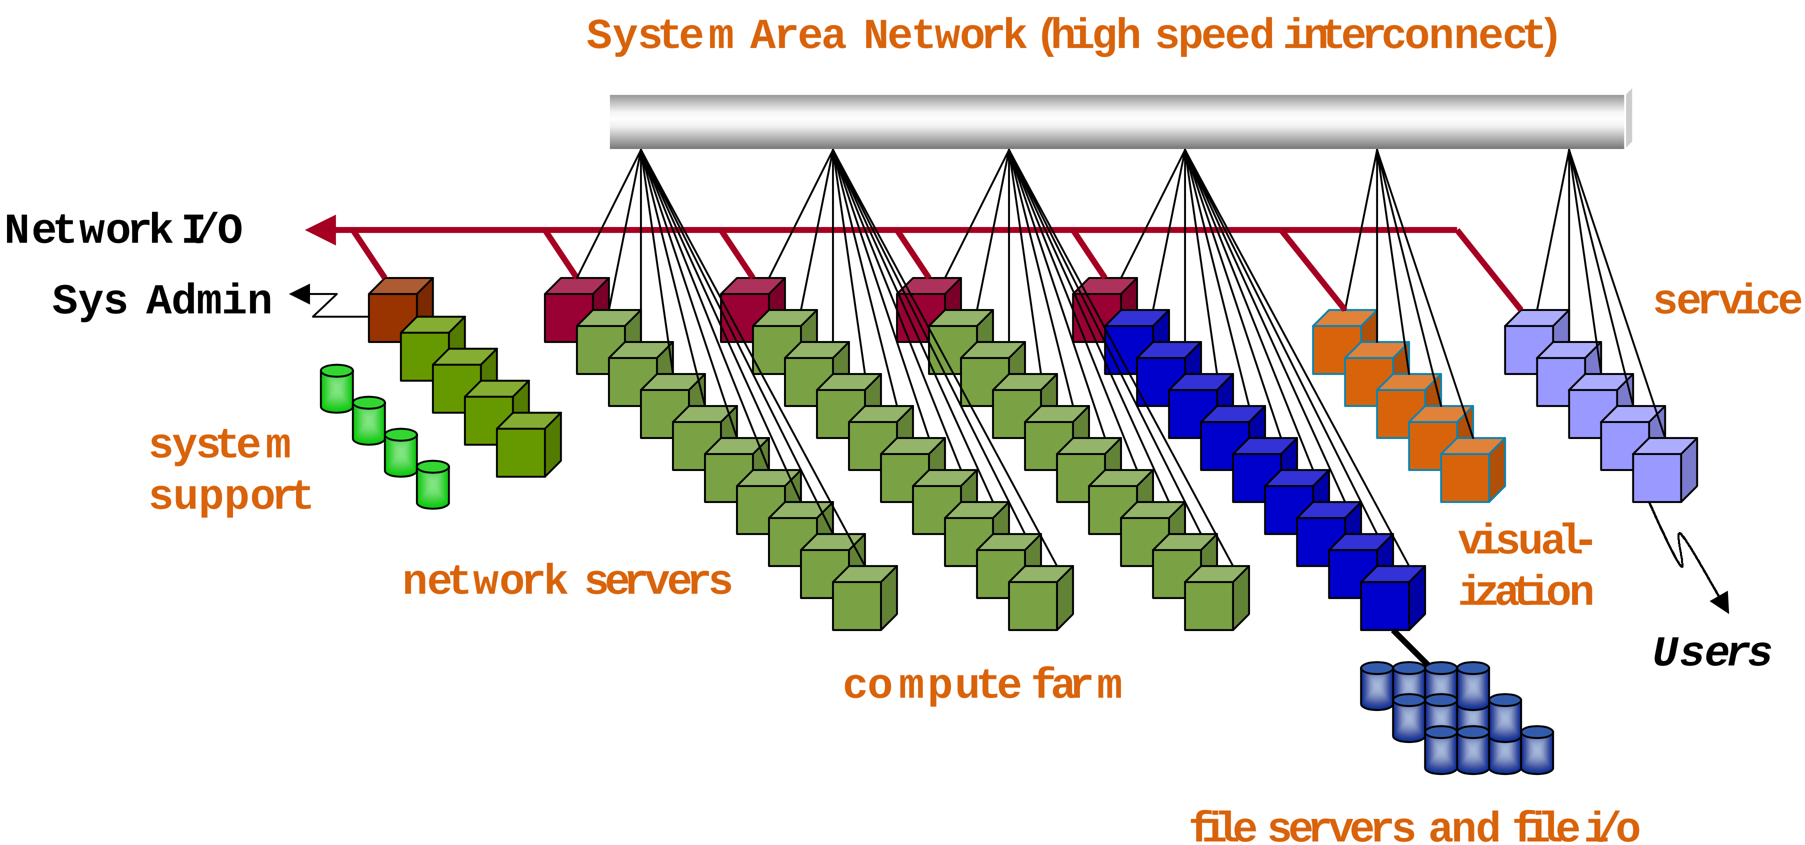

The Unanswered Questions

Cluster? SMP? Number of CPUs? 32 bit or 64 bit?

Memory? Disk? I/O? File System? Visualization?

Interconnect? Latency? Bandwidth? Cost of Ownership?

OS? Tools? Compilers? Schedulers? Debuggers?

ONLY WITH AN UNDERSTANDING OF APPLICATION(S) CAN ONE TRULY DESIGN AN APPROPRIATE SOLUTION.

HPTC Applications
(‘idealized’ technical requirements)

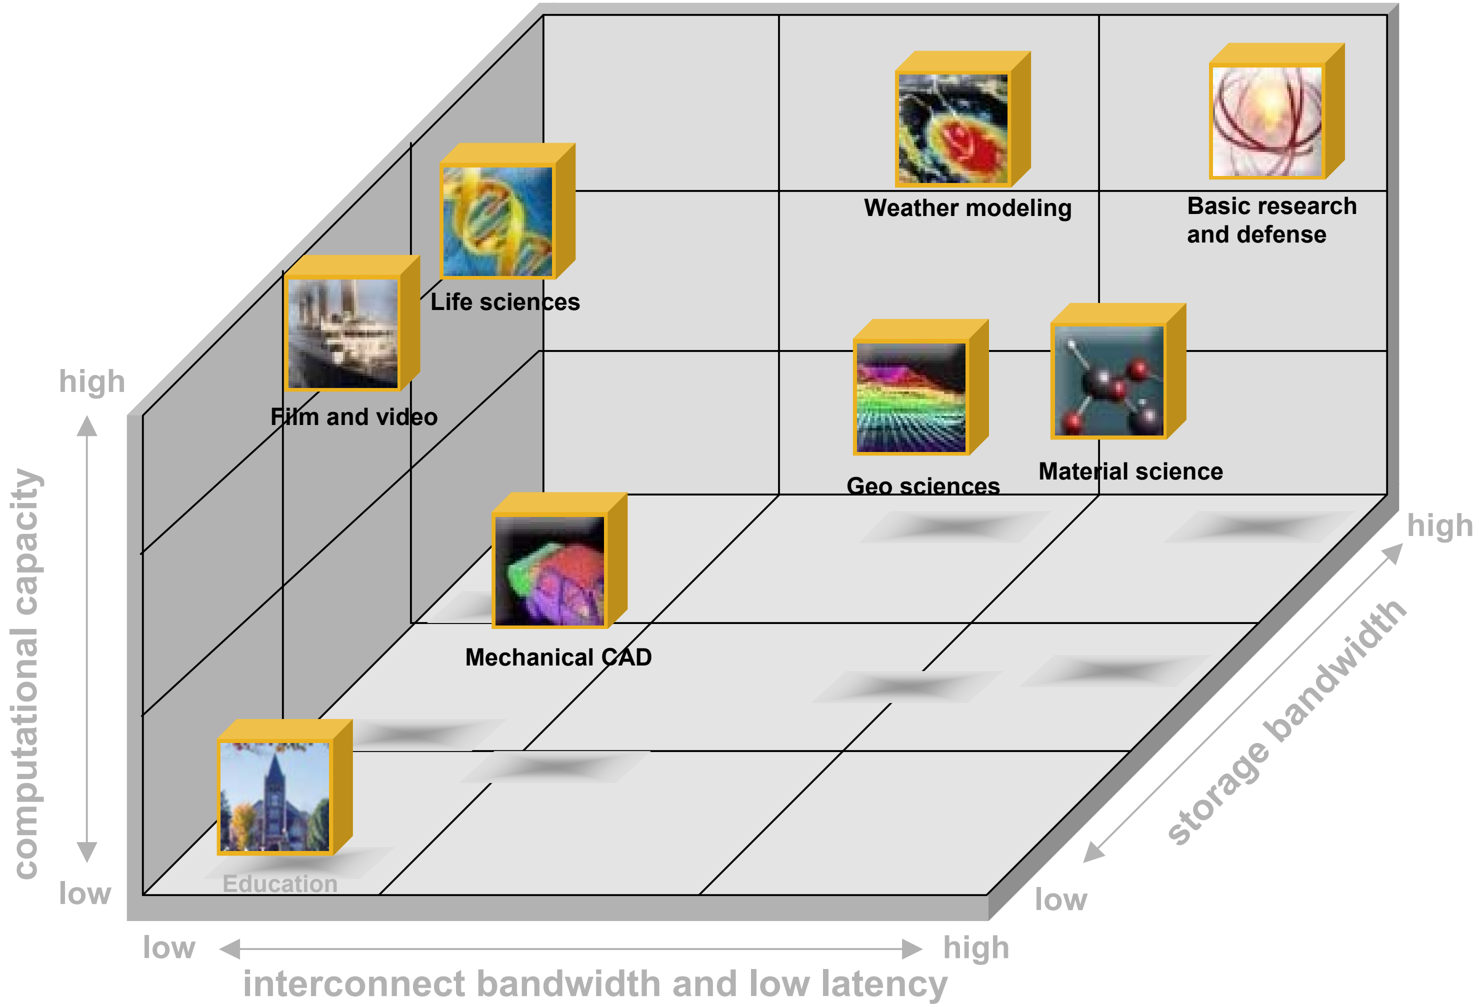

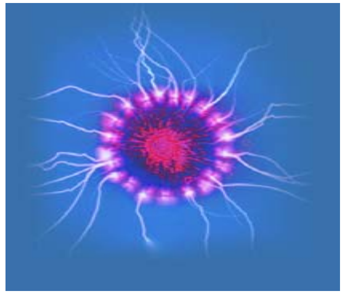

Life Sciences

Dave Mullally

HP WORLD 2003
Solutions and Technology Conference & Expo

titutions

hp's corporate clients include 11 of Fortune's top 13 Global 500 pharmaceutical companies

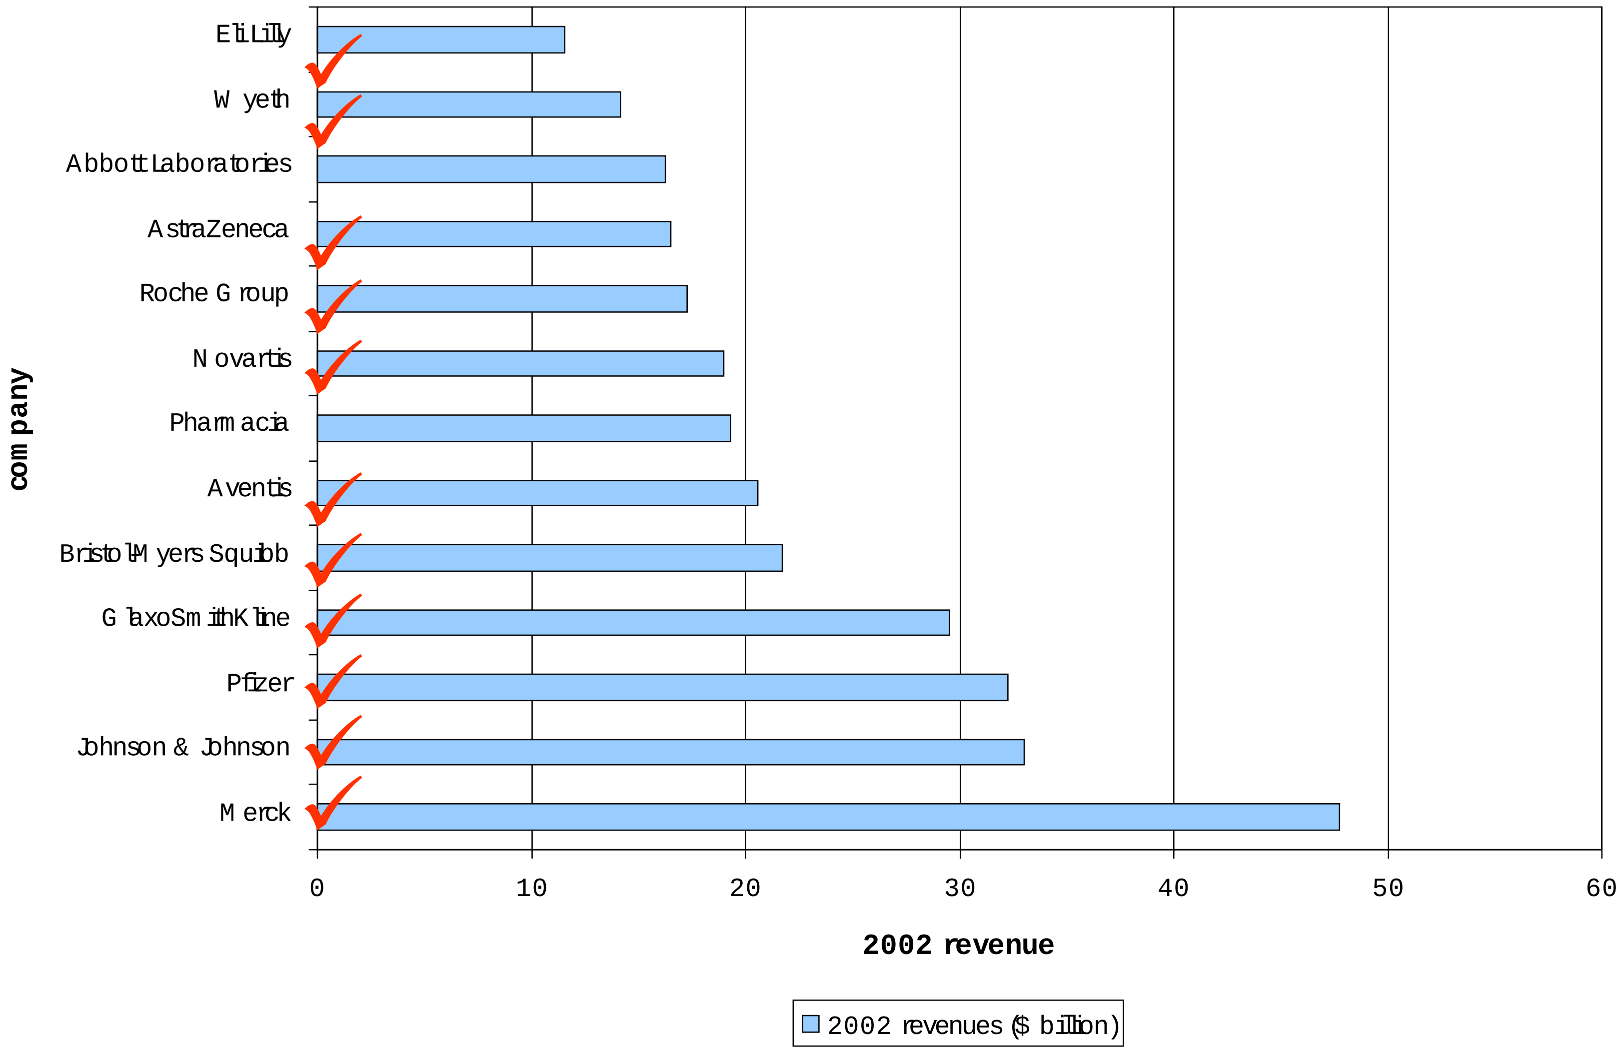

Source: Fortune Magazine, July 22 2002

Life Sciences Market Reach

HP WORLD 2003
Solutions and Technology Conference & Expo

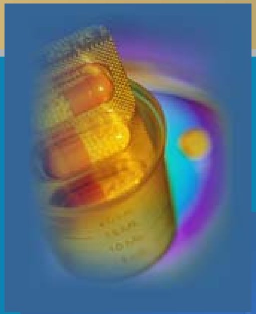

Pharmaceutical
- Drug discovery
- Pharmacogenomics
- Target gene therapy
- Aging
- Viral vectors

Medical
- Disease management
- Regenerative biology
- Cloning
- Diagnostics
- Disease research
- Reproduction

Agribusiness
- Hybrid processing
- Genetically designed
- Disease resistant strains
- Medical research
- Drug discovery
- Genetically modified stock
- Gene therapy
- Nutriceuticals

Bioinformatics

Genomics/Proteomics/ to System Biology

Computational Chemistry

Cheminformatics

Clinical Trials

Government
- Military
- Materials Research
- Environment clean-up
- Research/R&D
- Epidemiology

Environment
- Waste management
- Environment clean-up
- Detoxification

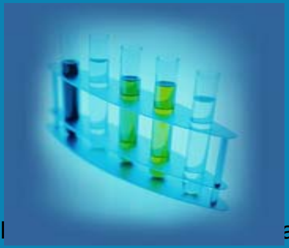

Chemical
- Environmental
- Manufacturing
- Petroleum
- Auto/Aerospace
- Personal Care
- Product

Academic Institutions
•Research

Life Sciences Customers and Market Sub-segments

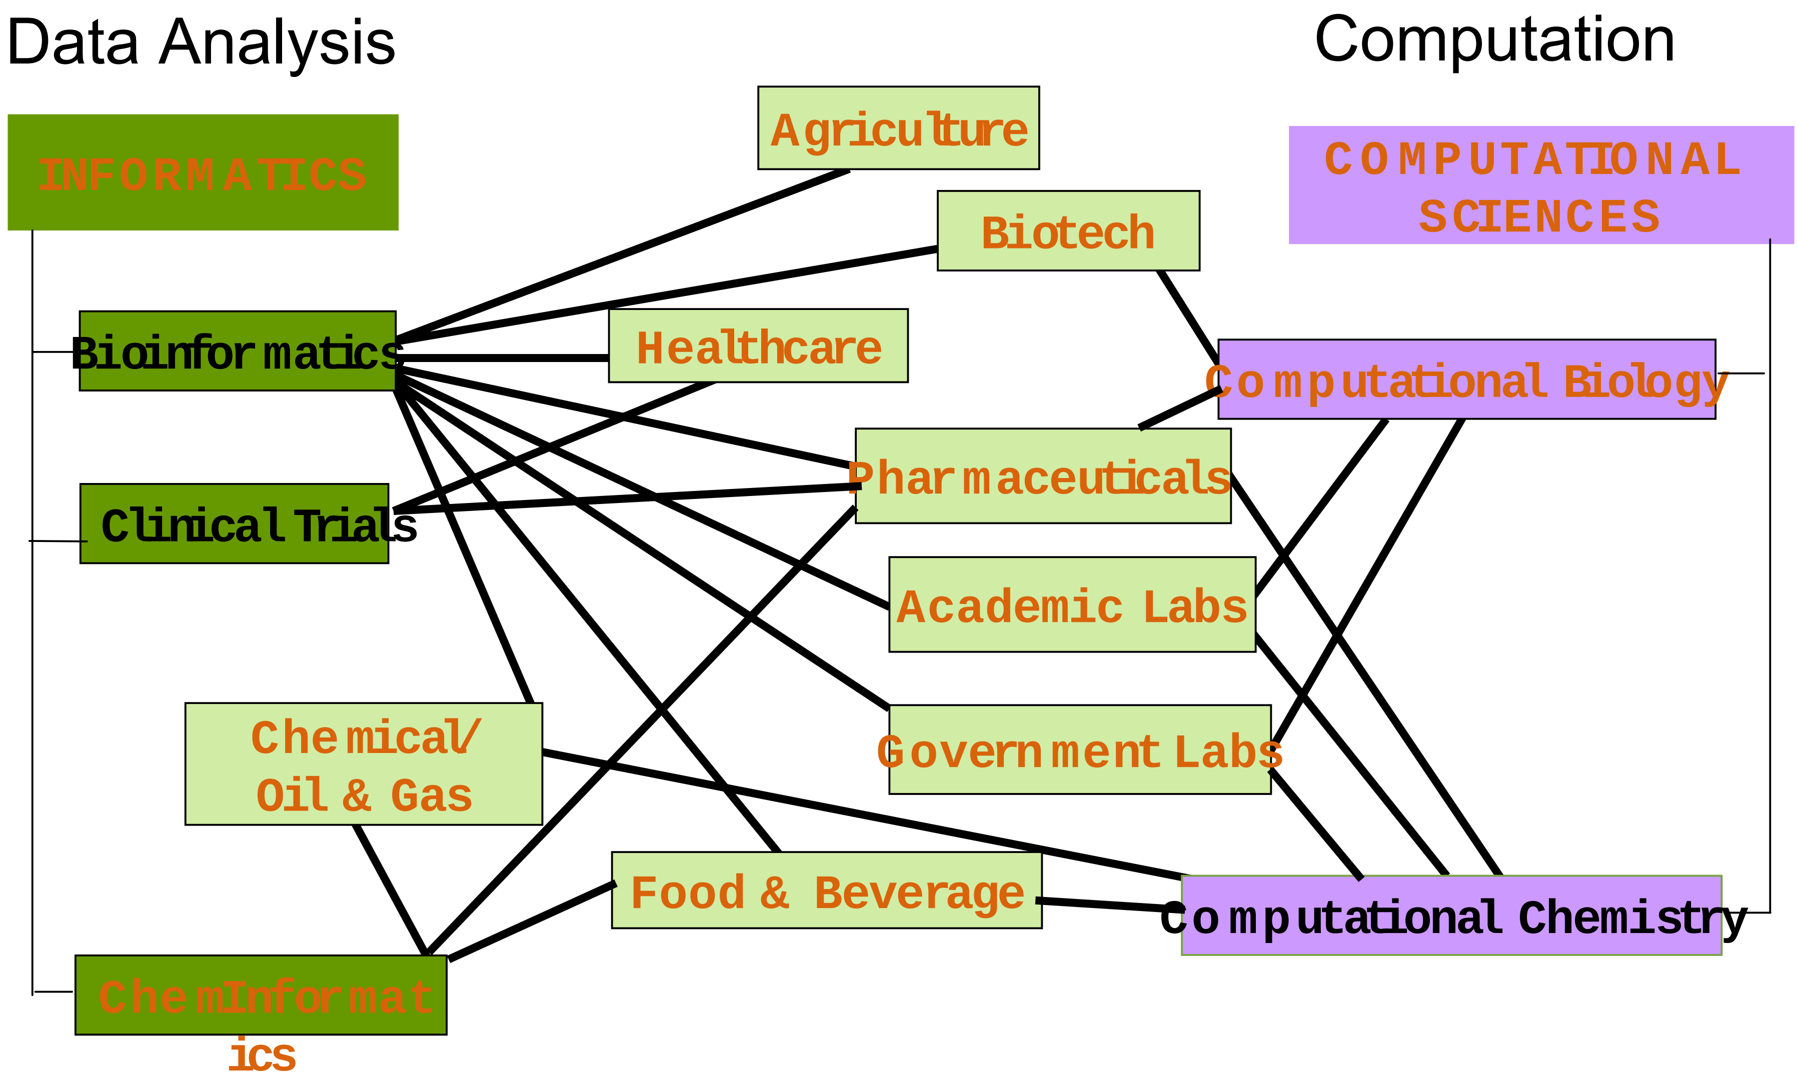

Life Sciences Computational needs

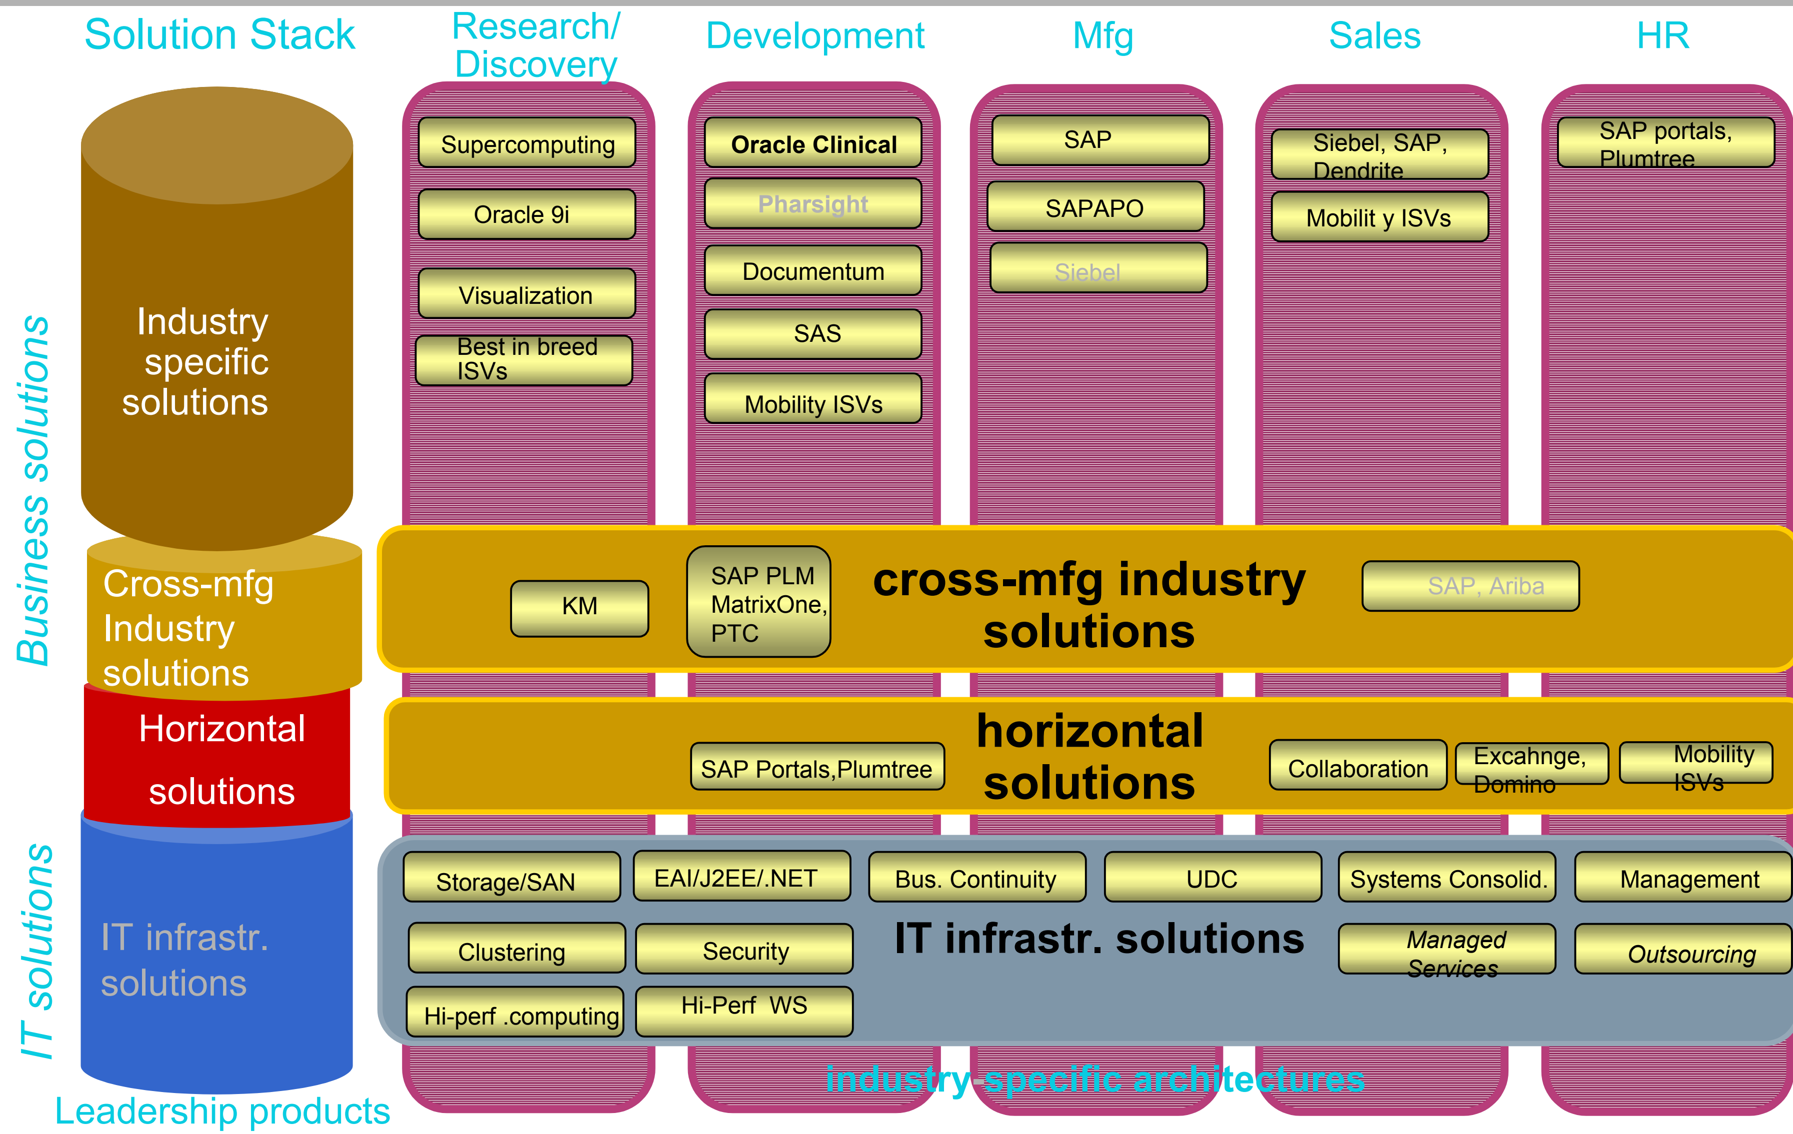

Hardware solutions for Life Sciences Customers

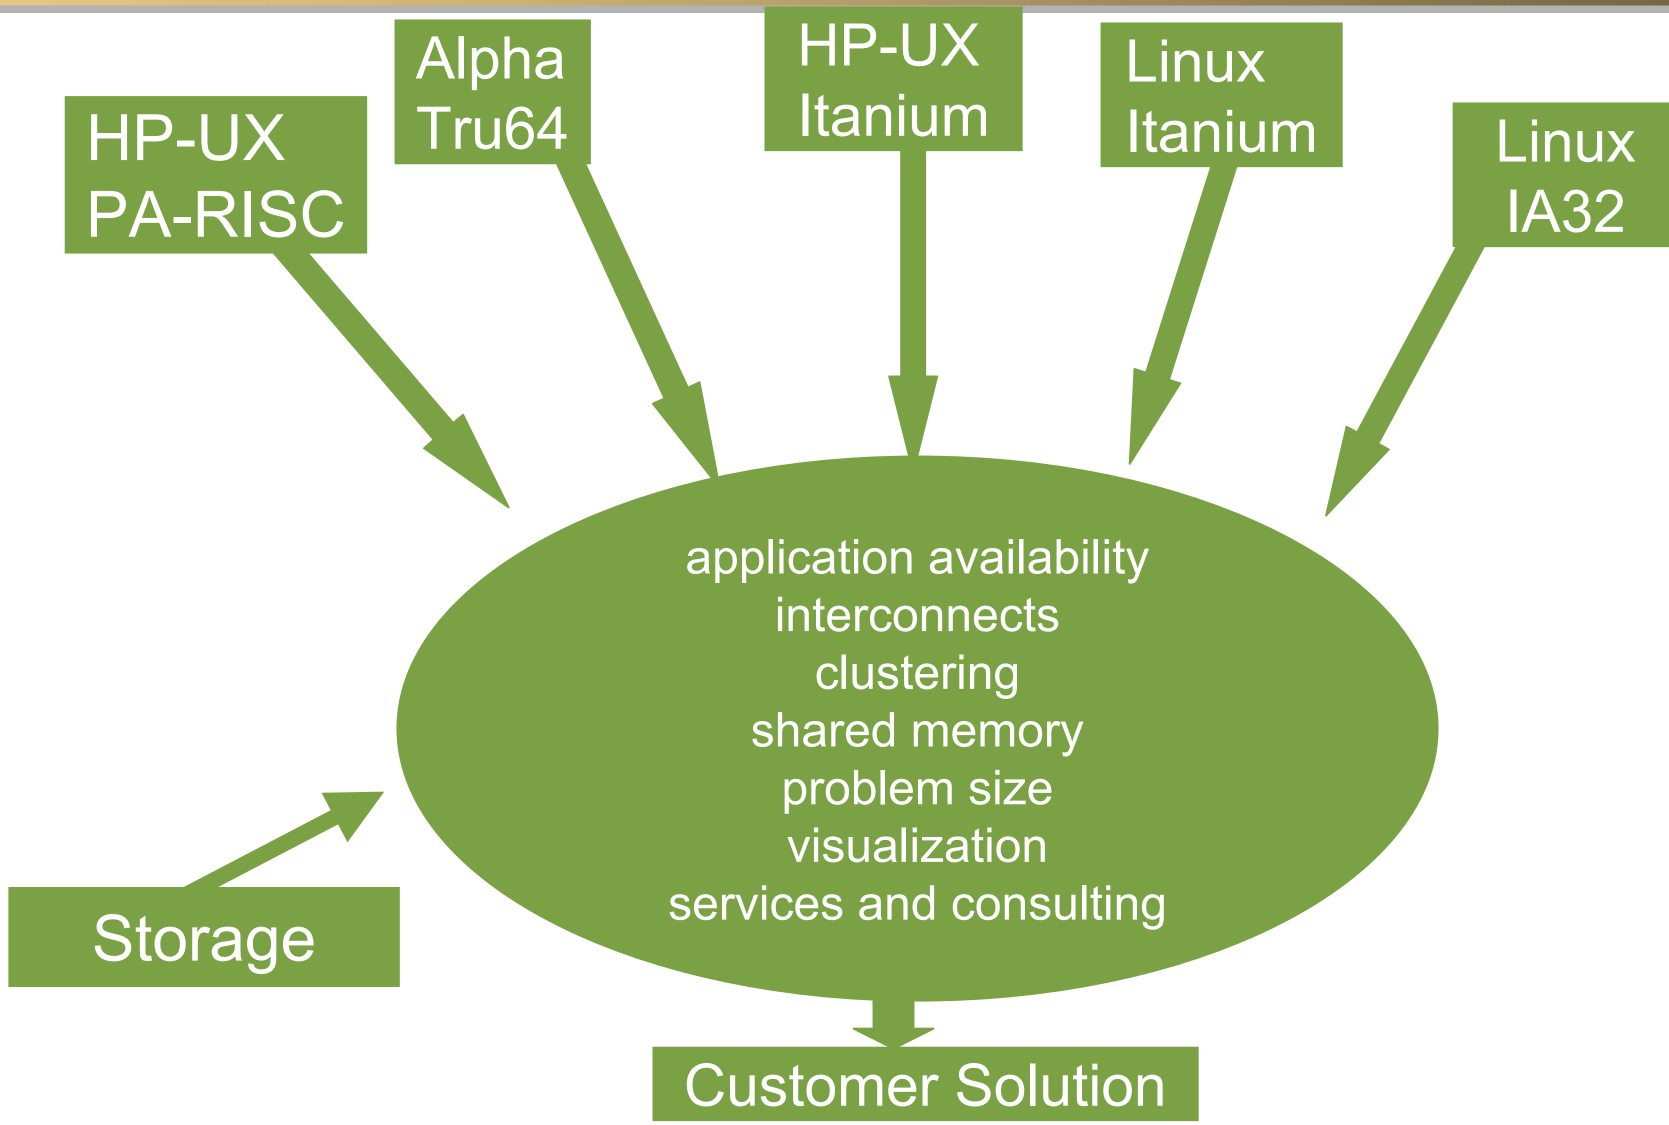

Application Support

- Ongoing marketing programs and technical support for current applications portfolio
- Continued collaboration on Alpha/Tru64U and PA-RISC/HPUX
- Long term strategy on a unified UNIX (HP-UX) and Linux portfolio for Itanium tm
- Specialized ISV Itanium tm developer programs “designing for the future”

Actual company and product names mentioned herein are the trademarks of their respective owners.

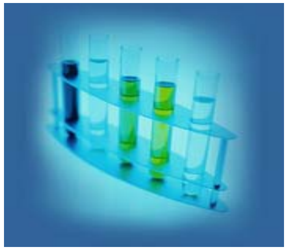

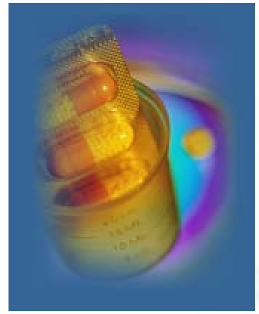

Computational Chemistry – issues and solutions

HP WORLD 2003
Solutions and Technology Conference & Expo

There is no Single Typical Computational Chemistry System

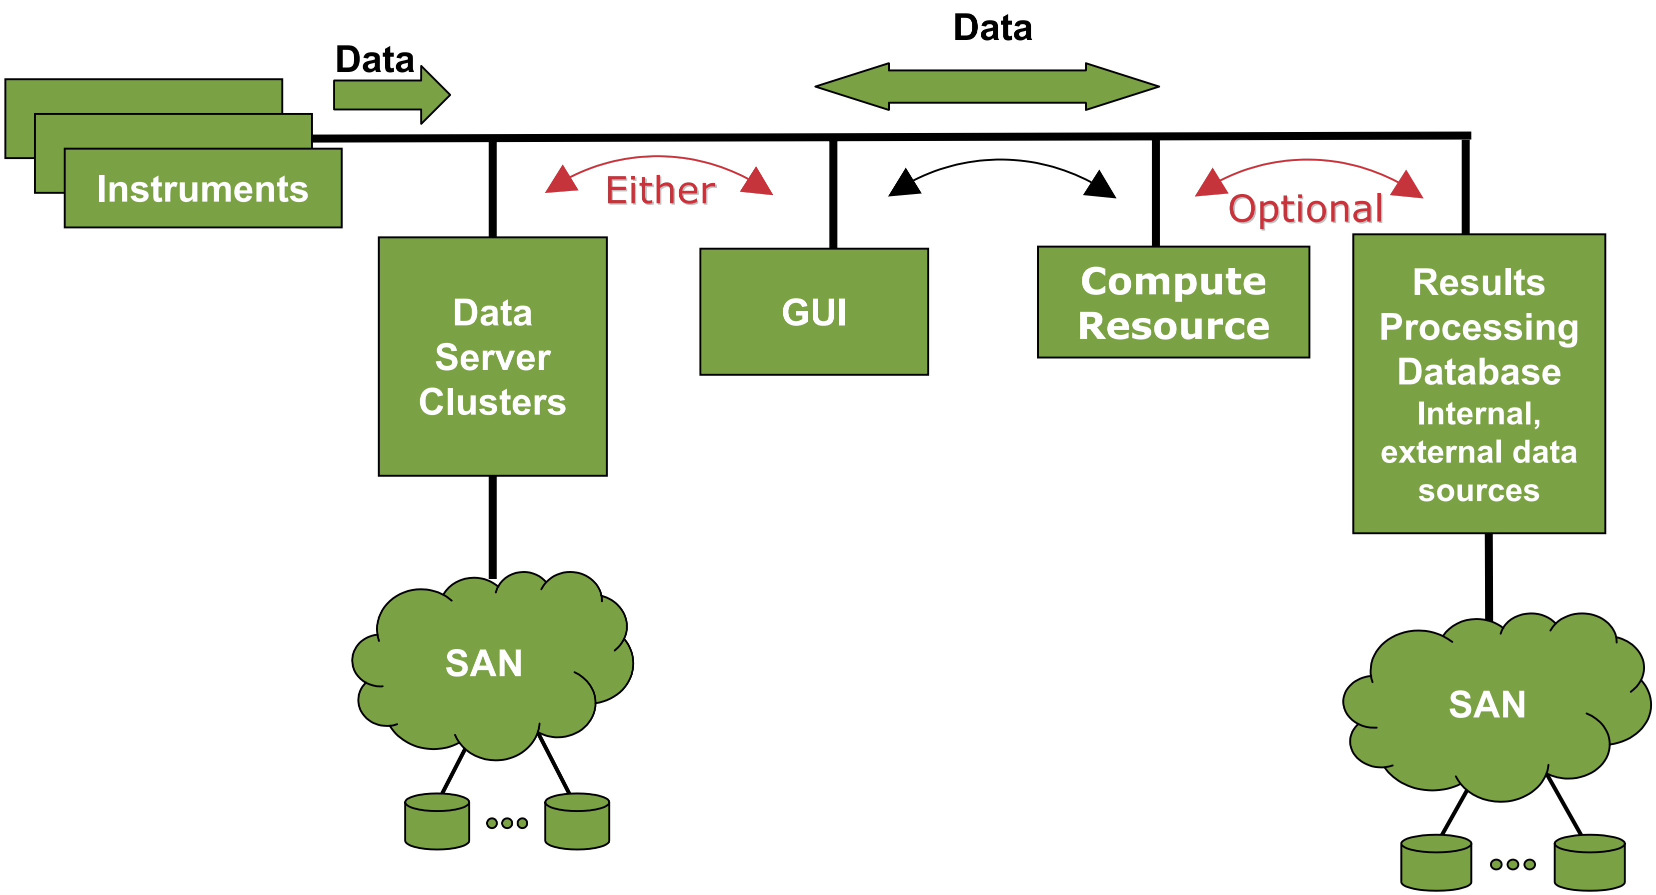

What is Driving IT in Computational Chemistry?

- Users tend to be groups of 3-10 scientists
  - Customers typically experts
- Quantum Chemistry
  - Used on small rigid molecules (like drugs) where the energy of bonding is important
  - Hardware requirements
    - Large memory – Cache Size, 64-bit
    - Large files – I/O, 64-bit
    - Parallelism - SMP
- Molecular Dynamics
  - Used on large floppy molecules where how the molecule folds is important
  - Hardware requirements
    - Scatter/Gather – Cache Size, Memory Latency
    - Parallelism – Divide and Conquer, Inter-connect

Status of Key Computational Chemistry Applications

ISV	APPLICATION	IPF HP-UX	IPF LINUX	PA HP-UX	ALPHA TRU64	IA-32 LINUX	Memory Bandwidth Cache Size I/O Inter-Connect 64-bit
Gaussian	Gaussian03	Soon	X	Soon	X	X	MB CS I/O Linda 64-bit
Tripos	SYBYL, Unity	Beta		X		Beta	MB CS
OpenEye	ROCS	X	Soon	X	X	X	MB
UCSF	AMBER	X	X	X	X	X	MB CS MPI (Div & Conq)
PNNL	NWChem	X	X	X	X	X	MB Global Arrays 64-bit
Iowa State	GAMESS	X	X	X	X	X	MB MPI I/O

Gaussian 03 – Parallel Scaling

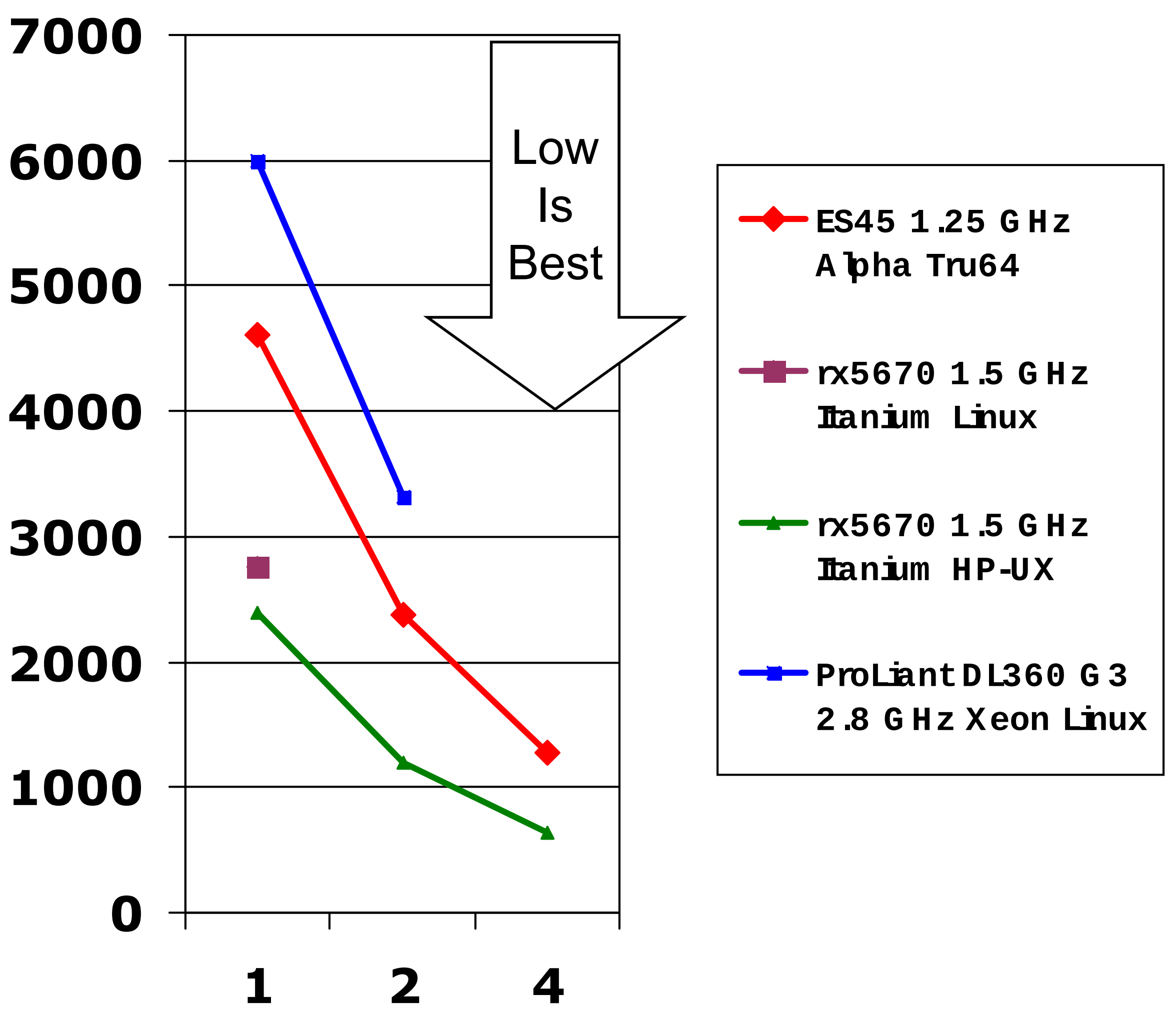

- **Alpha Pinene**
  C_1 Symmetry
  RB3LYP/6-31G*
  182 Basis Functions
  FREQUENCY

Quantum Chemistry Benchmark Exercises:

Memory Bandwidth

Cache Size

SMP Parallel good to 16 CPUs

GAMESS (UK) – Performance

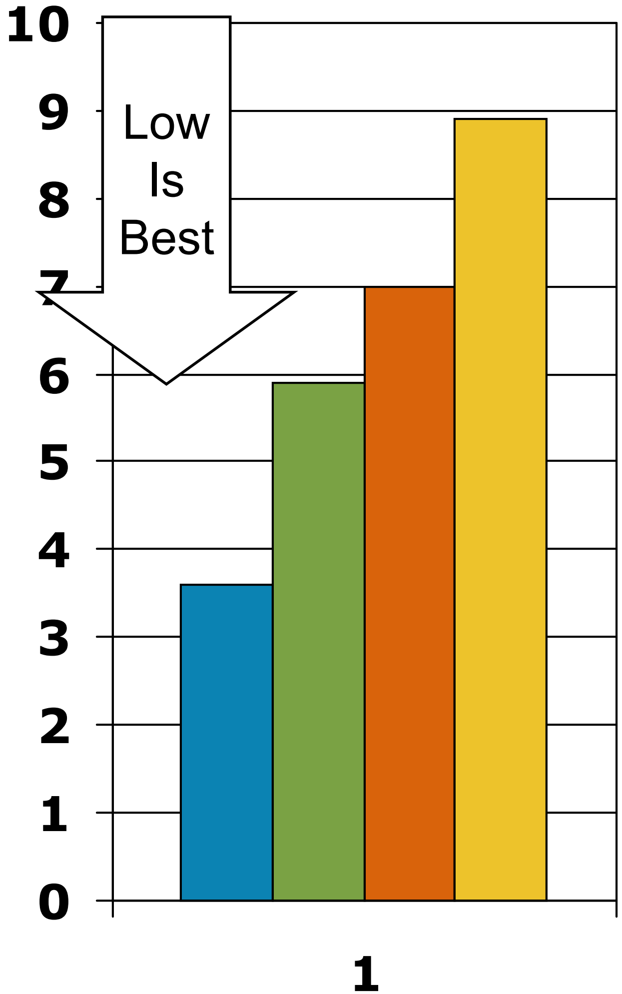

- rx5670 1.5 GHz Itanium HP-UX
- IBM-p690 1.7 GHz power4 AIX
- Pentium 4 2.666 GHz Linux
- AMD-Opteron 244 1.8 GHz Linux

- SCF Kernel

Quantum Chemistry Benchmark Exercises:

Memory Bandwidth

Cache Size

http://www.dl.ac.uk/CFS/benchmarks/compchem.html

AMBER – Parallel Scaling

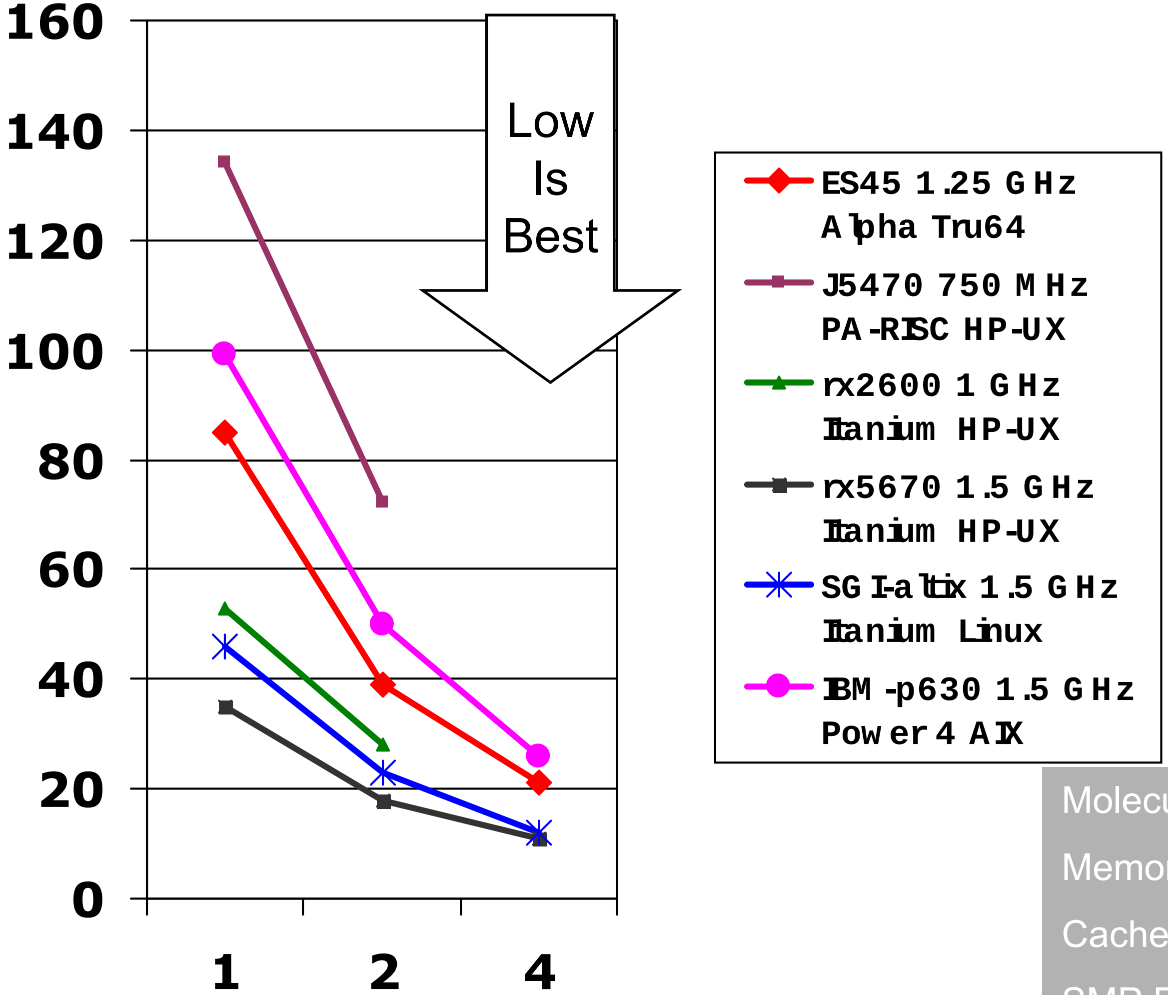

- **JAC Benchmark**

Number of Atoms = 23558

Number of Water molecules = 7023

Box size = 62.230 (X), 62.230 (Y), 62.230 (Z)

Cutoff = 9.00 A

Ewald Coefficient = 0.30768

Interpolation Order = 4

Time step = 1 fs

Molecular Dynamics Benchmark Exercises:

Memory Bandwidth

Cache Size

SMP Parallel

64-bit versus 32-bit

- 64-bit
  - Proprietary operating systems offer lowest cost of ownership
  - Some Quantum Chemistry methods require 64-bit
  - SIMD offers better parallel scaling than clusters
  - Madison gives best single processor performance today!
- 32-bit
  - Best price/performance
  - Best application availability

HP-UX versus Linux on Itanium

- HP-UX
  - More efficient compilers so typically better performance
  - More stability
  - OS scales, so better for SuperDome, better for SIMD applications
- Linux
  - Better at clustering (XC)
  - Open source

Summary of Computational Chemistry Customer Requirements

- Customers typically require 5-10 applications
- Architecture/operating varies with customer priorities
  - Applications often determine solution
  - Large shared memory users best solution is Itanium/HP-UX
  - 64-bit cluster users best solution is Itanium/Linux
  - 32-bit clusters best price performance

rx2600

SuperDome

DL380 G3

rx5670

Blade Servers

ProLiant Clusters

Itanium-based clusters

XCxxxx

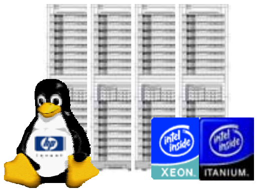

(mid 2003)

AlphaServer GS

AlphaServer SC supercomputer

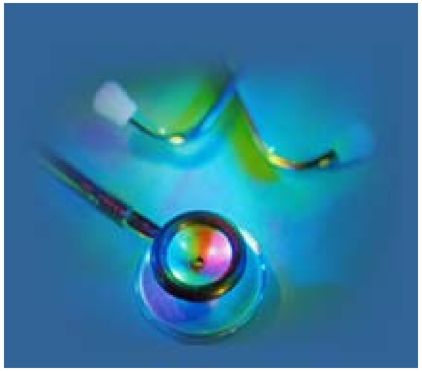

Bioinformatics – issues and solutions

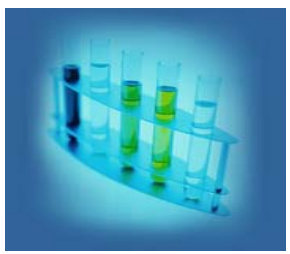

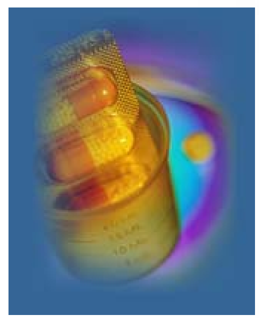

HP WORLD 2003
Solutions and Technology Conference & Expo

Typical Bioinformatics System

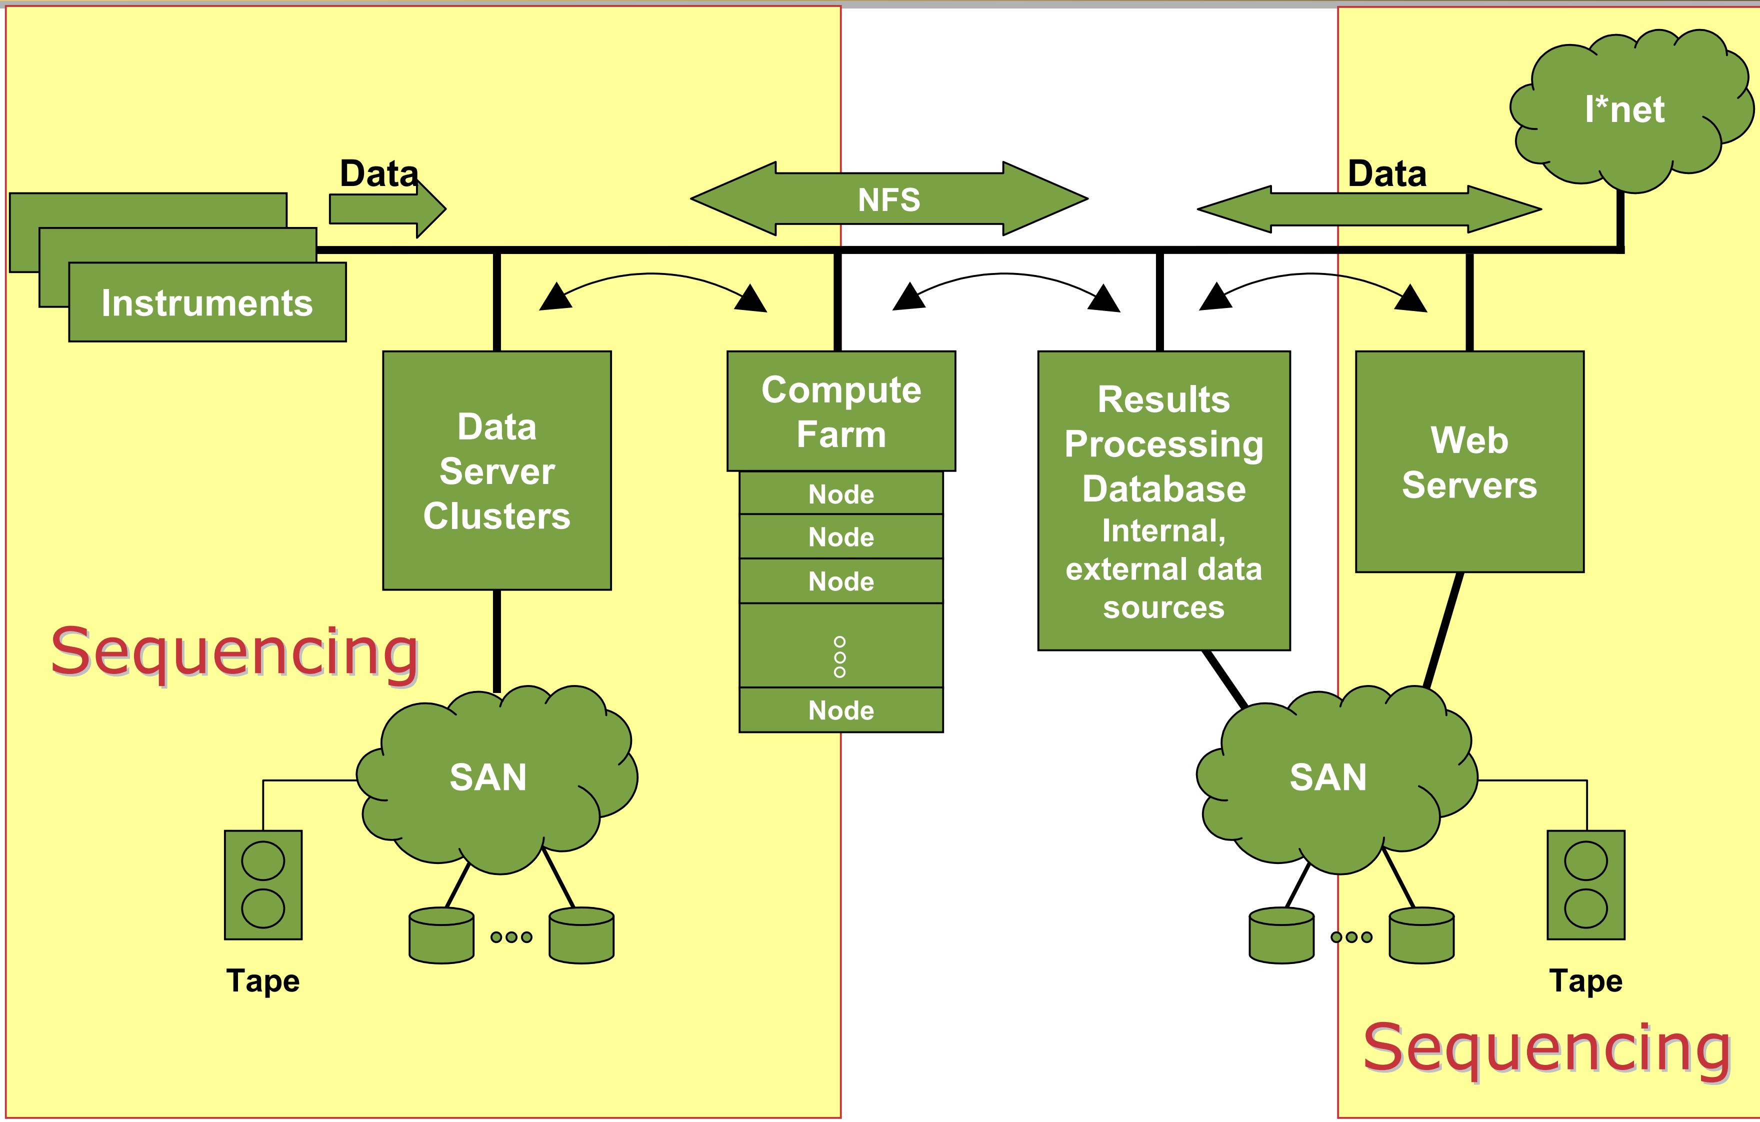

What is the IT direction in Bioinformatics?

- Rapid development phase of applications
  - Numerous new applications
  - Most applications are academic
  - Home grown coded prevalent
  - Linux is dominating
- Divide and Conquer methods viable
  - Parallelism on clusters is excellent – Inter-connect important
- Explosion in data
  - Large storage critical
  - 32-bit is adequate for Sequencing
  - 64-bit is required for analysis of large databases in Proteomics/Genomics
    - Large memory – Memory Bandwidth

Status of Key Bioinformatics Applications

ISV	APPLICATION	IPF HP-UX	IPF LINUX	PA HP-UX	ALPHA TRU64	IA-32 LINUX	Memory Bandwidth Cache Size I/O Inter-Connect 64-bit
NCBI	BLAST	X	X	X	X	X	MB CS Div & Con
Accelrys	GCG				X		MB CS
Wash. U	BLAST	X	X	X	X	X	MB CS Div & Con
U Virginia	FASTA	X	X	X	X	X	MB CS MPI
LION	SRS, bioScout	target	target		X	soon	MB CS 64-bit
Geospiza	RepeatMasker	X		X	X		MB CS
Wash. U	HMMER	X	X	X	X	X	MB CS

Blast Performance

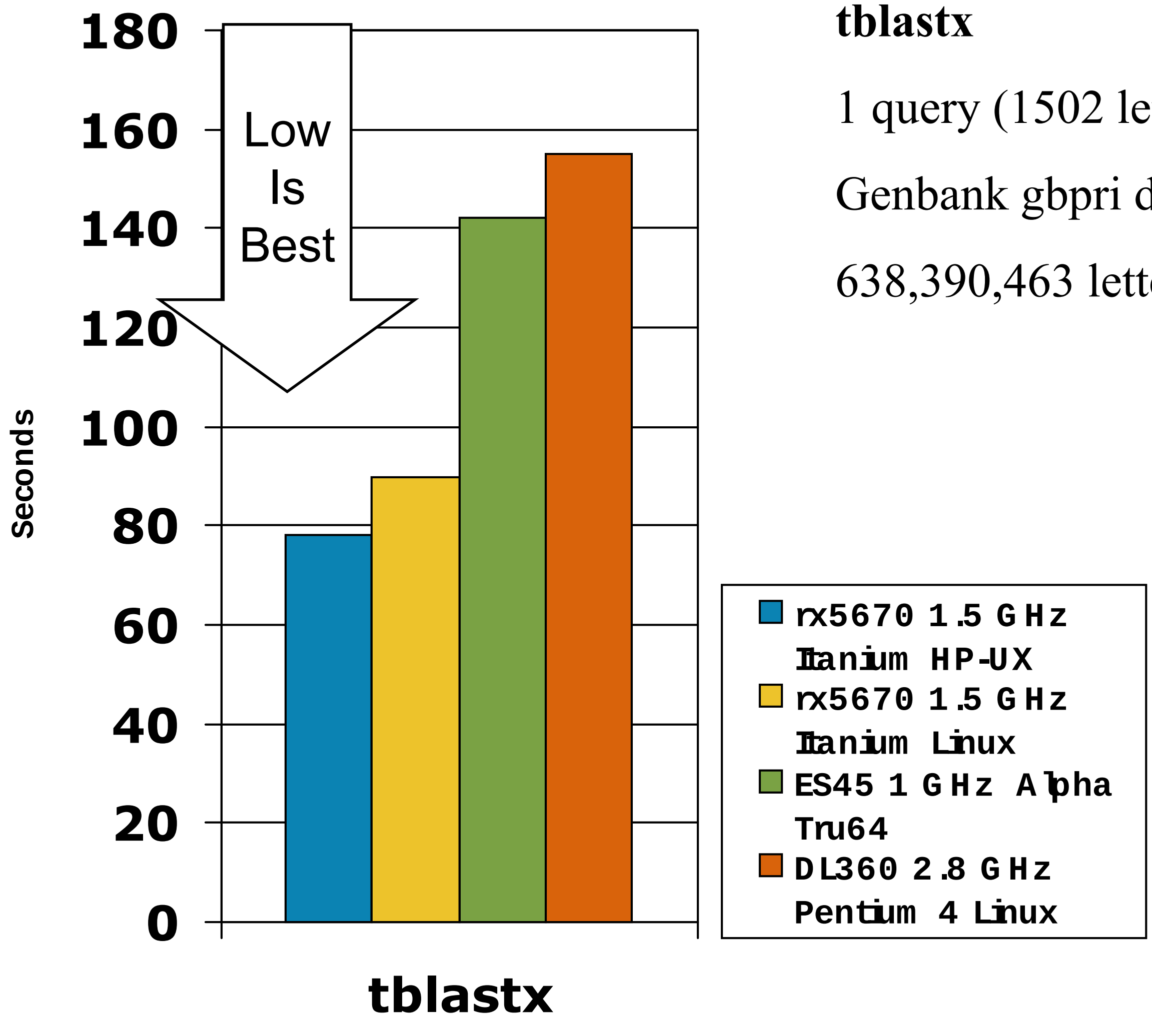

tblastx

1 query (1502 letters)

Genbank gbpri database with 109,867 sequences

638,390,463 letters

Sequencing Benchmark Exercises:

Memory Bandwidth

Cache Size

Summary of Bioinformatics Customer Requirements

- Customers typically require 10-20 applications
- Storage is Critical
- Linux dominates
  - Sequencing customers buy 32-bit clusters
  - Proteomics/Genomics customers buy 64-bit clusters
  - Alpha has significant installed base

Itanium-based clusters

XCxxxx

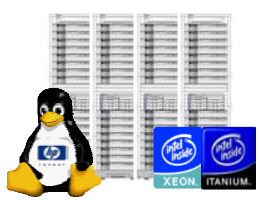

(mid 2003)

AlphaServer GS

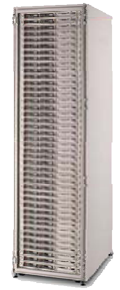

ProLiant Clusters

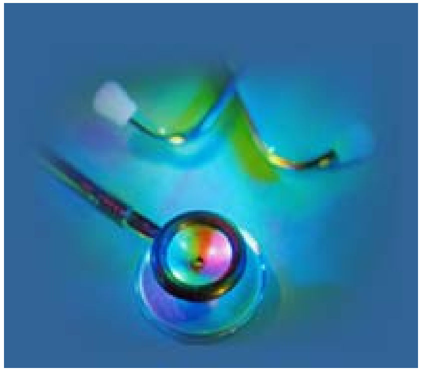

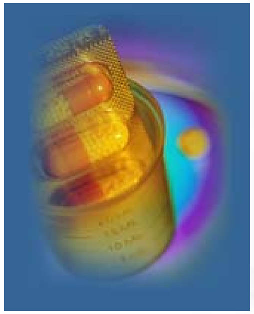

Clinical Trials – issues and solutions

HP WORLD 2003
Solutions and Technology Conference & Expo

Typical Clinical Trials System

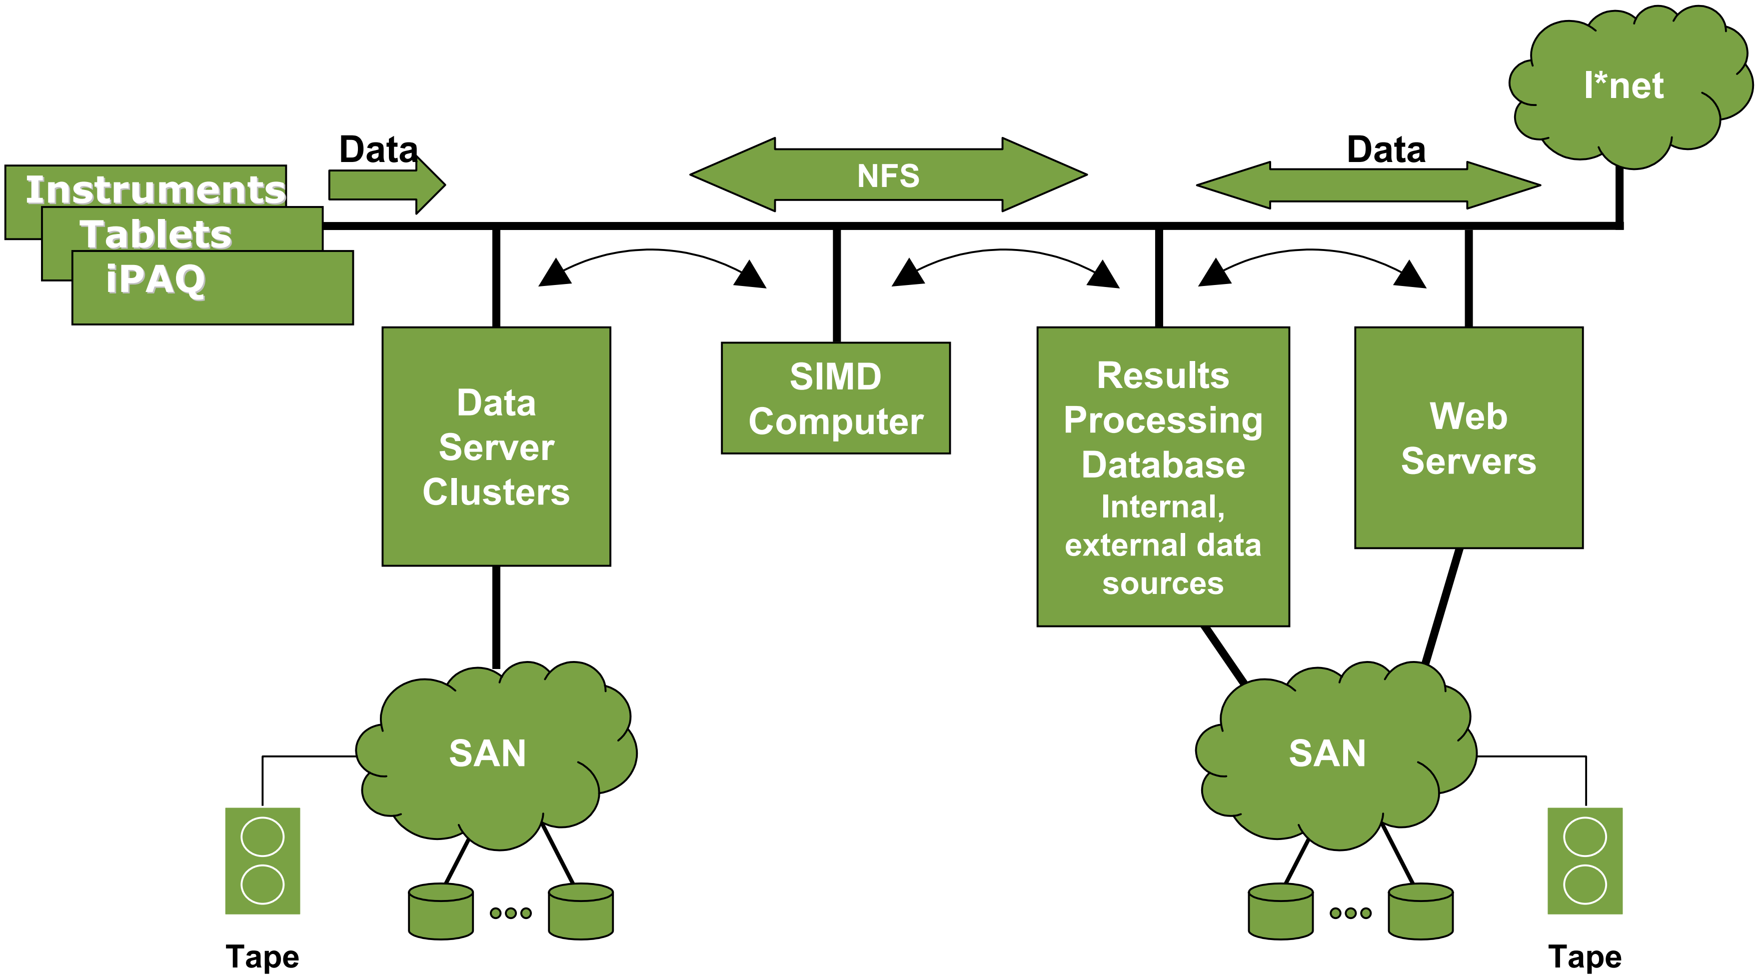

What is driving IT in Clinical Trials?

- Data collection is creating the data explosion
  - Storage and management important
- Need for consulting
  - End to end clinical trials process is supported by a fabric of non-collaborating applications, systems, and data bases
  - Mergers & acquisition magnify the challenges
  - The system is supported by data managers in a labor intensive process
- The process MUST meet all regulatory requirements, is quality assurance compliant, and facilitates audit requirements

Status of Key Clinical Trials Solutions

Issue	IPF HP-UX	IPF LINUX	PA-RISC HP-UX	ALPHA TRU64	IA-32 LINUX	Memory Bandwidth Cache Size I/O Inter-Connect 64-bit
FDA Certification	X		X	X		
Oracle Clinical	FY04		X	X		CS I/O
Phase Forward	FY04					CS I/O
SAS	FY04		X	X		CS I/O
eResearch Technology			X			CS I/O

Performance – Transactions Per Minute

Computer	Processor	# cpu	tpm/cpu
Superdome	875 MHz PA-RISC	64	6615
ES45	1.25 GHz Alpha	4	14094
Superdome	1.5 GHz Itanium	64	11048
rx5670	1.5 GHz Itanium	4	21935
HP ProLiant DL380-G3	1.4GHz Pentium III	1	18051

Clinical Trials Data Conversion Challenge

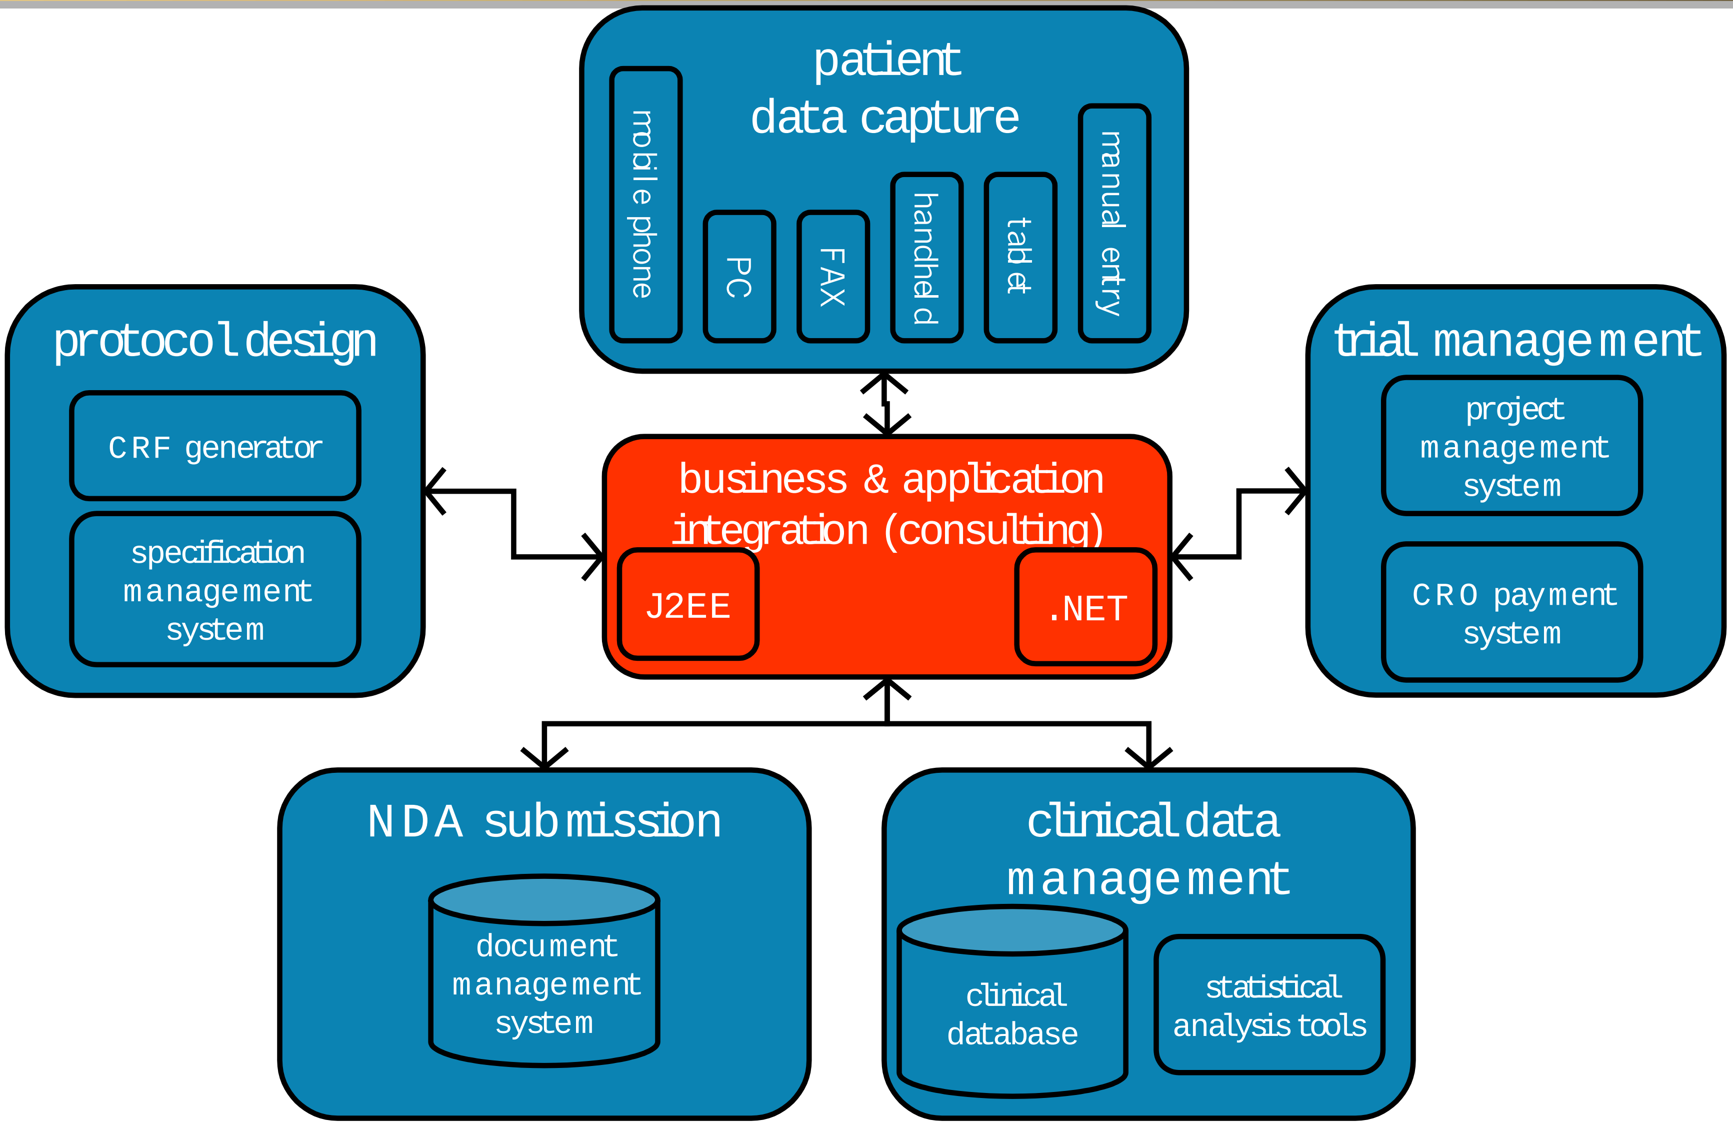

Summary of Clinical Trials Customer Requirements

- Customers typically require 2-3 applications
- FDA requirement means
  - Linux is not viable
  - Customers buy 64-bit Alpha and PA-RISC solutions
- Large shared memory machines providing most tpm, so are preferable to clusters
- Consulting is critical

AlphaServer
SC supercomputer

SuperDome

Conclusion

ONLY WITH AN UNDERSTANDING OF APPLICATION(S) CAN ONE TRULY DESIGN AN APPROPRIATE SOLUTION

Computer Aided Engineering

Curtis Bennett

Solution Architect

Americas HPTC Solution Design Team

HP WORLD 2003

Solutions and Technology Conference & Expo

What is CAE

Computer Aided Engineering

CAE is a broad term used for the use of computers to design, analyze, and manufacture products and processes.

Why use CAE

Harvard Business Review
Enlightened Experimentation: The New Imperative for Innovation by Stefan Thomke

"Experimentation lies at the heart of every company's ability to innovate. In other words, the systematic testing of ideas is what enables companies to create and refine their products."

Having enough computer resources to quickly test different ideas earlier in the design cycle will provide better and more innovative products (DOE, Stochastic Ananalysis).

What is CAE used for

- Simulation Based Design
  - virtual testing during the component design process
  - virtual trial and error of design concepts
  - optimization of component performance, quality and manufacturability
- Virtual Product Development
  - Save time and cost with realistic virtual testing
  - Improves quality by merging product design and process engineering
  - Leverages creativity and innovation
- Virtual Prototyping
  - product testing in realistic conditions
  - appropriate design decisions
- Virtual Manufacturing
  - process engineering under realistic conditions
  - evaluation and choice among various process solutions.

CAE Process

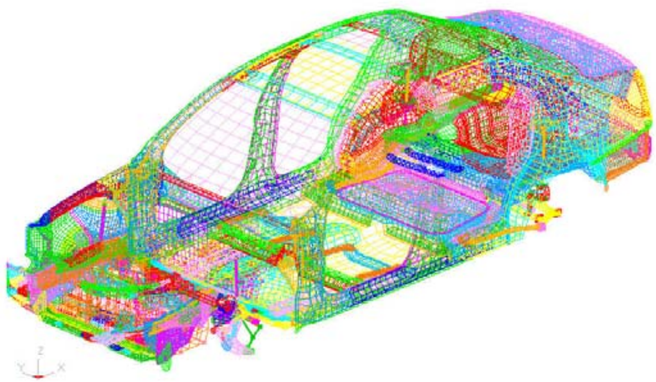

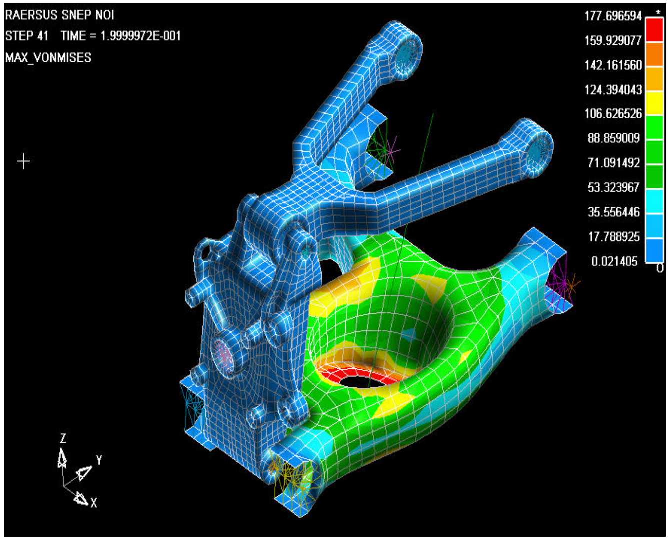

Pre-Processing

CAD geometry is used to create computational meshes

Post-Processing

The results from the analysis are reviewed to assess product performance

Solution

Loading and Boundary conditions are prescribed for the analysis and the job is submitted to a compute server

Application Areas

- Structures
  - Implicit structural analysis, such as statics, dynamics, and thermal analysis.
  - Explicit structural analysis focuses on transient dynamics and quasi-static analyses using an explicit approach appropriate in many applications such as drop test, crushing and many manufacturing processes.
- Computational Fluid Dynamics
  - A complex set of algorithms used for modeling and simulation of fluid, gas, heat or material flows.
- Electromagnetics
  - Electromagnetic simulation can be applied to a broad range of industries including transducer design, rotating machines, induction heating process simulation, ion optics and RLC parameter extraction for electronic engineering.

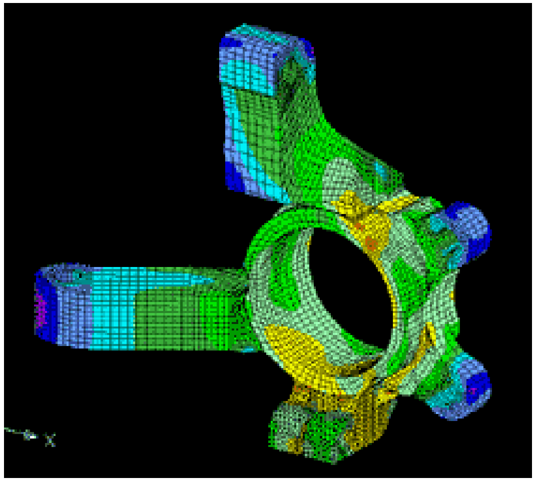

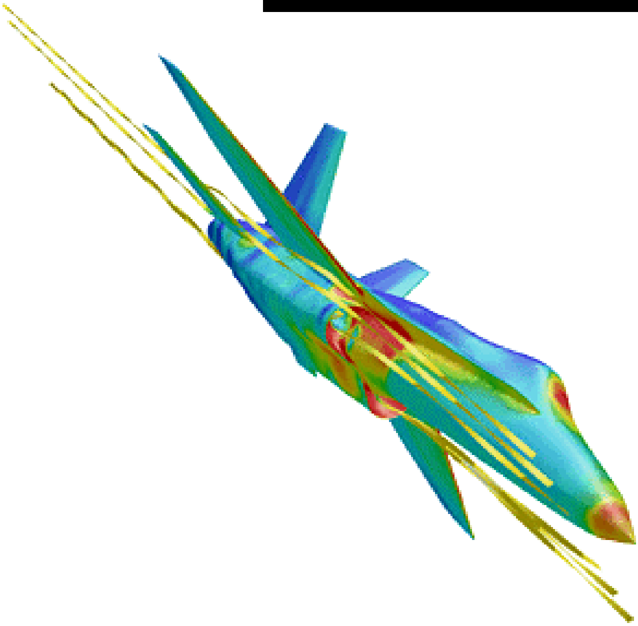

Where is Structural Analysis Done

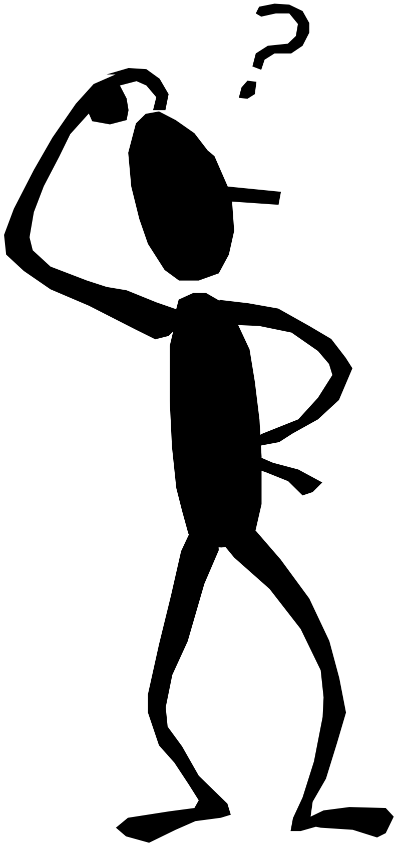

Automotive

Crash Analysis

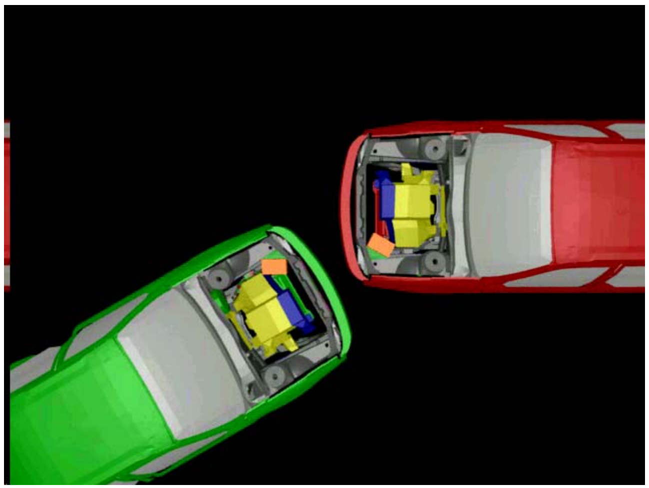

Airbag Deployment

Component Analysis

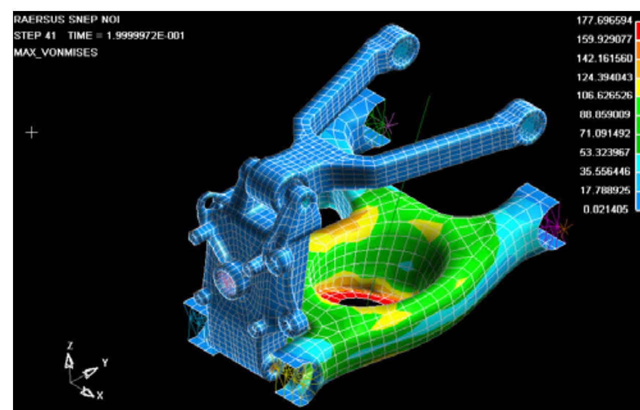

Occupant Dynamics

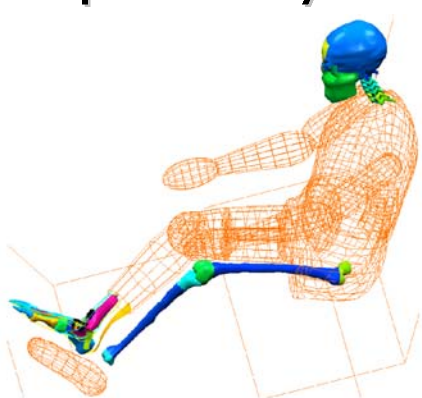

Aerospace and Defense

Landing Gear

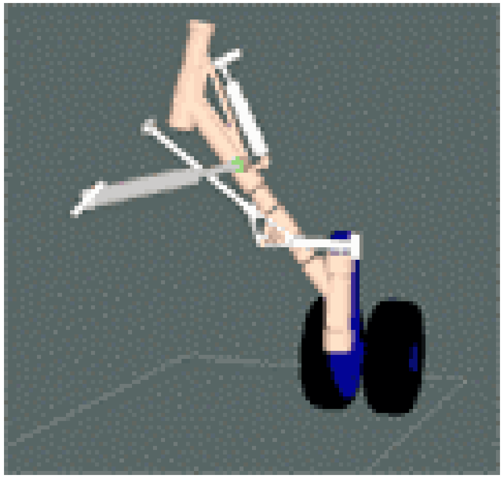

Failure Modes

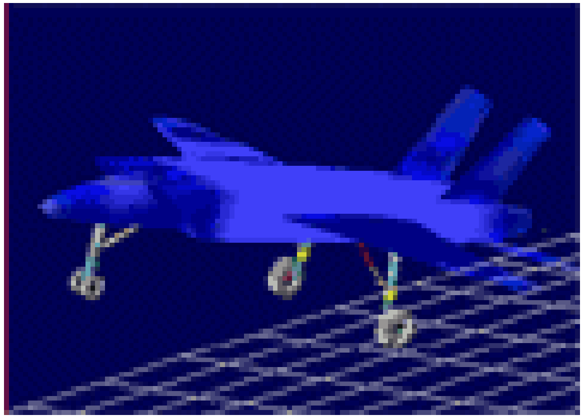

Ground Vehicle Durability

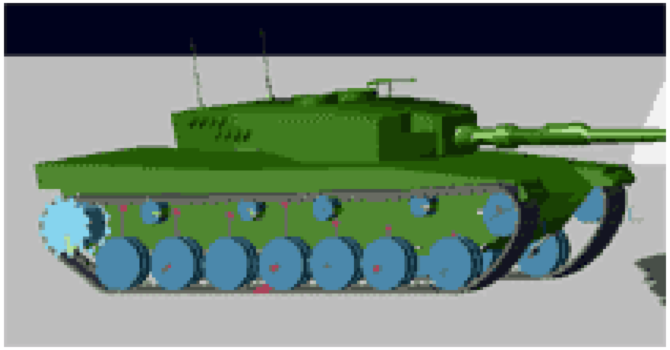

Manufacturing Process

Metal Forming

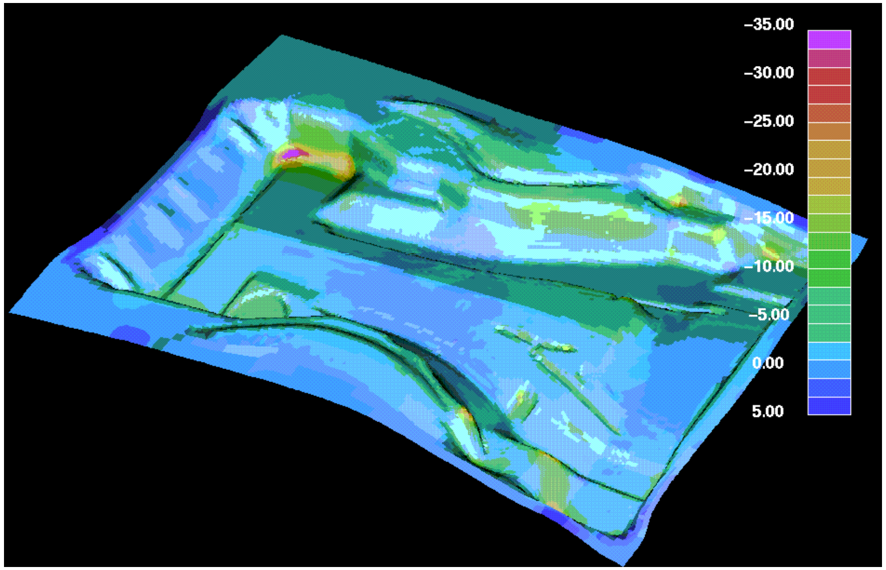

Injection Molding

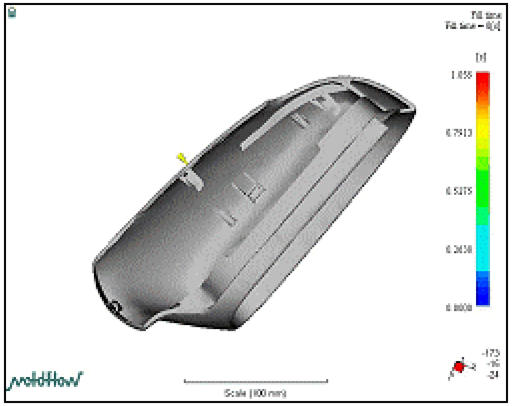

Biomechanics

Orthopedic Implants

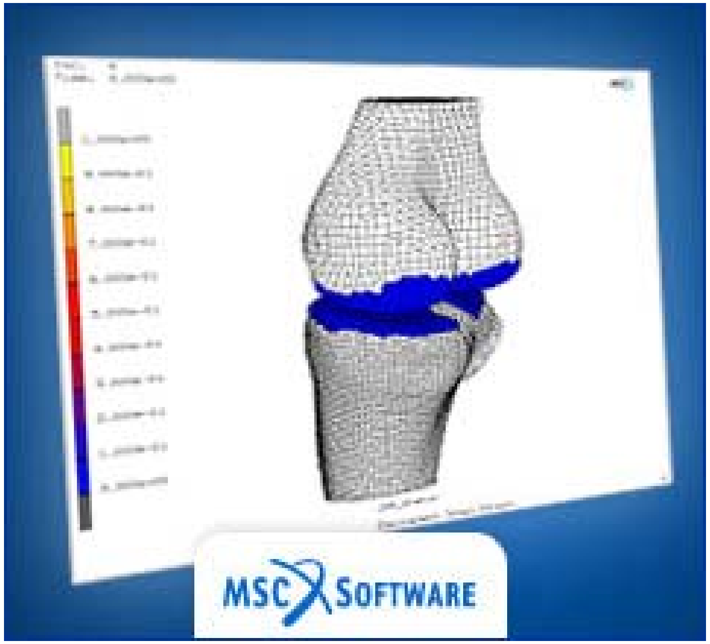

Oral Implants

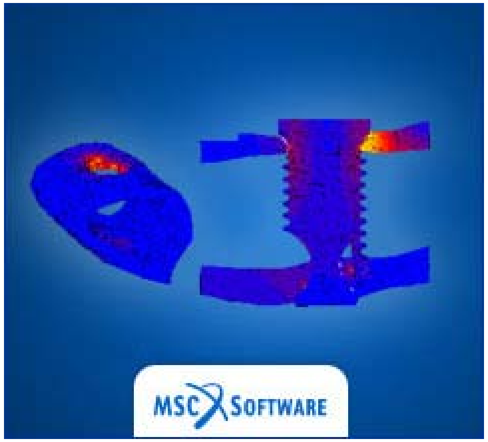

Consumer Products

Clothing

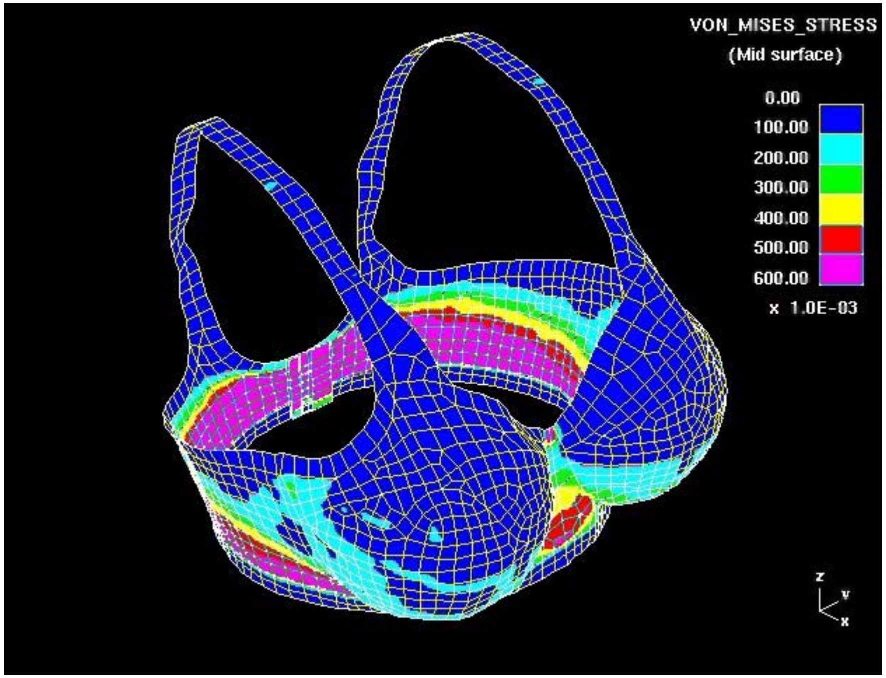

Washing Machine

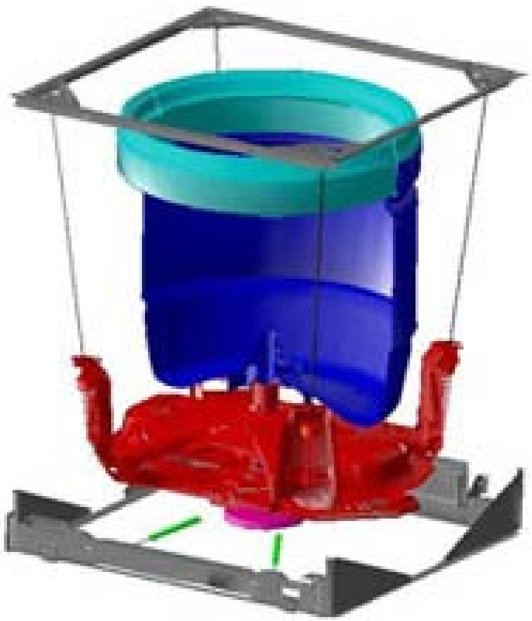

Where is CFD Used

Automotive

Climate Control

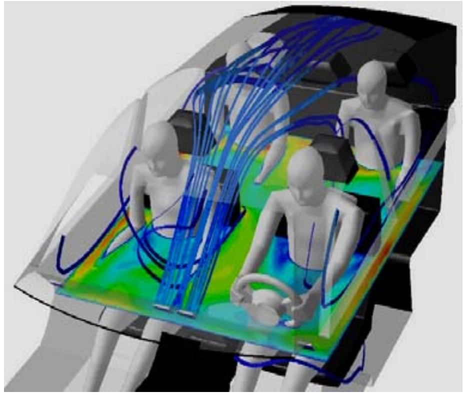

Vehicle Aerodynamics

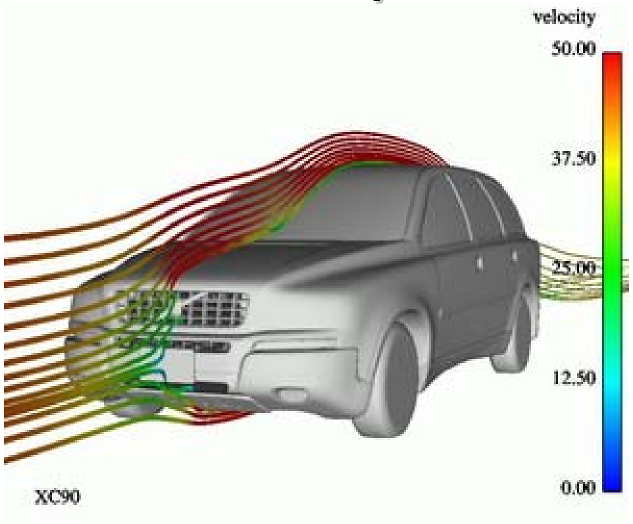

Thermal Management

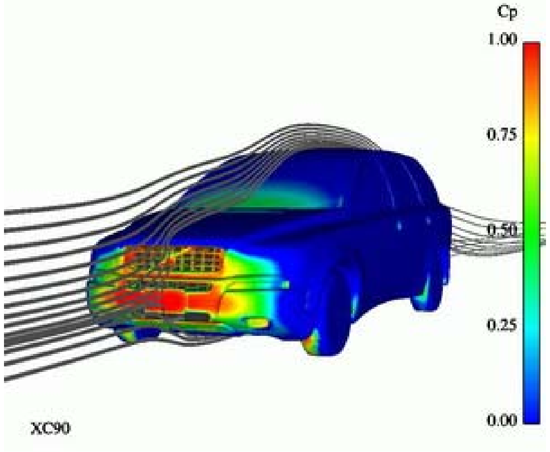

Internal Flows

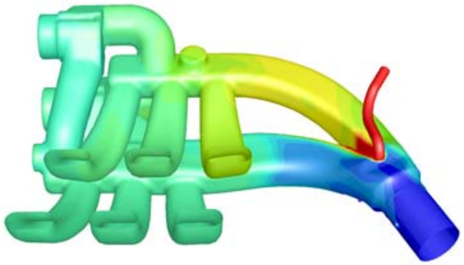

Aerospace and Defense

Load Separation

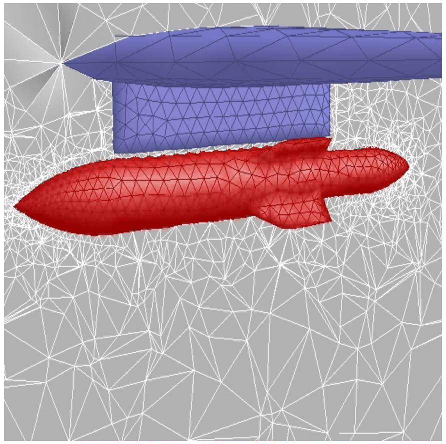

Helicopter Flow Field

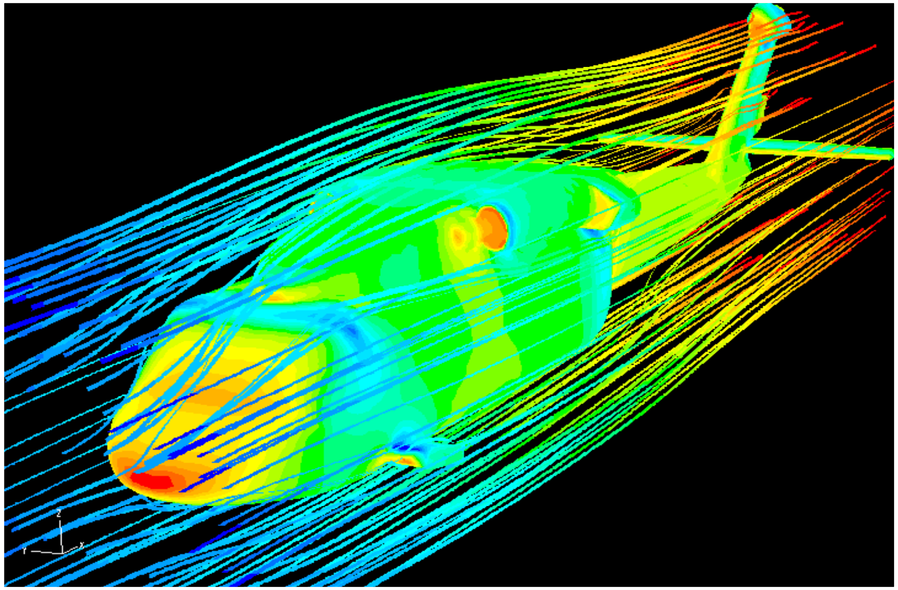

Fin Optimization

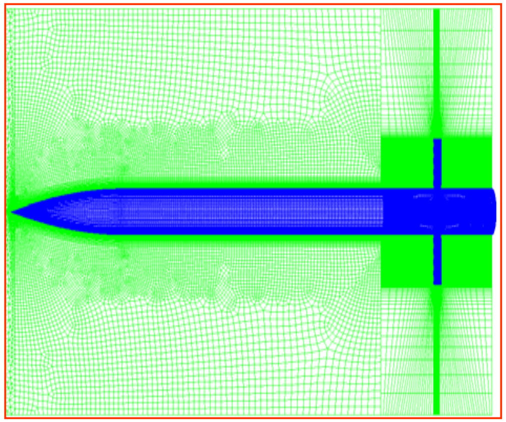

Biomedical

Flow in an Artery

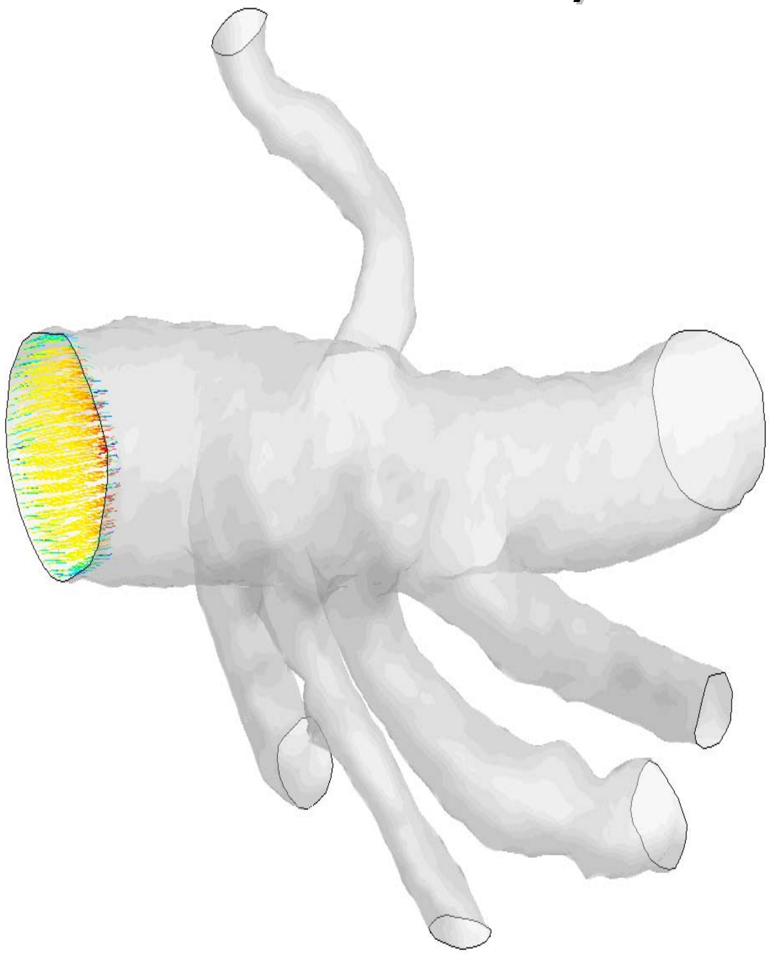

Chemical Processing

Oil and Water Separation

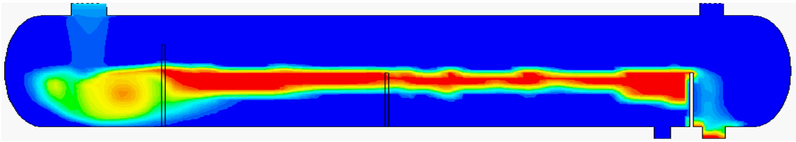

Archimedes Screw

Tank Filling

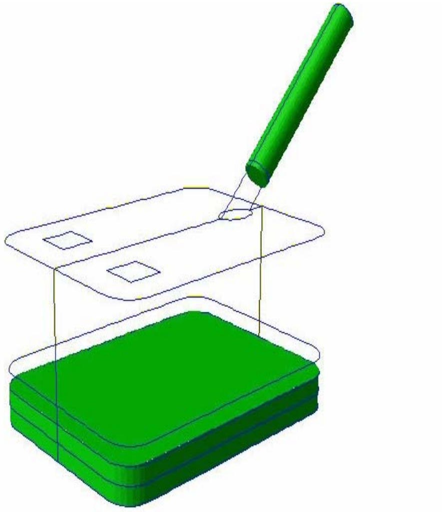

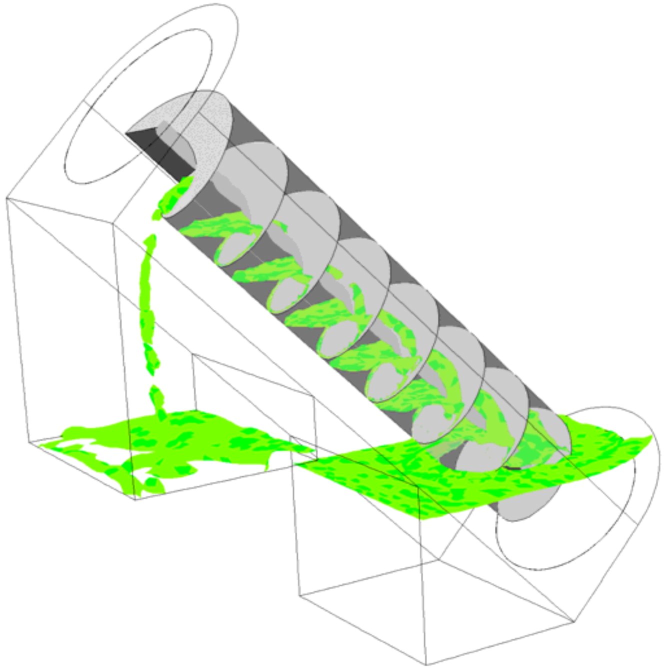

Power Generation

Turbine Combustor

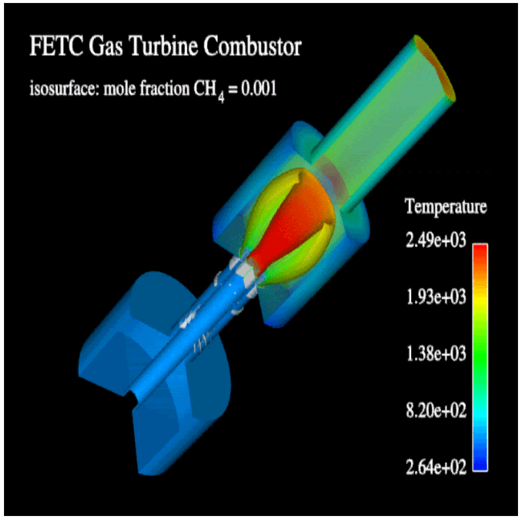

Fuel Cell

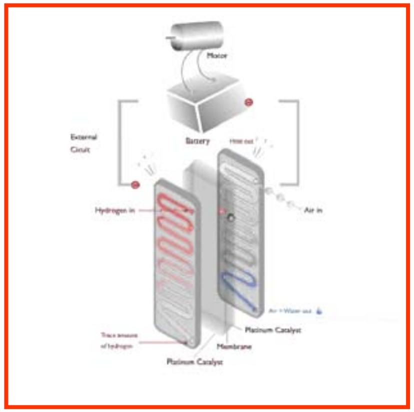

Nuclear

Manufacturing Process

Polymer Extrusion

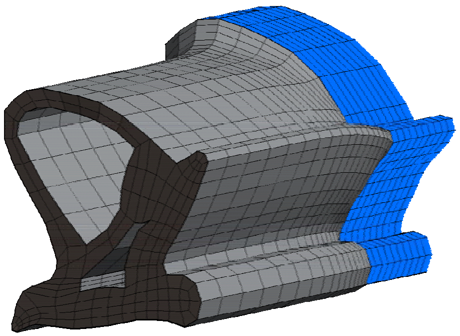

Glass Furnace

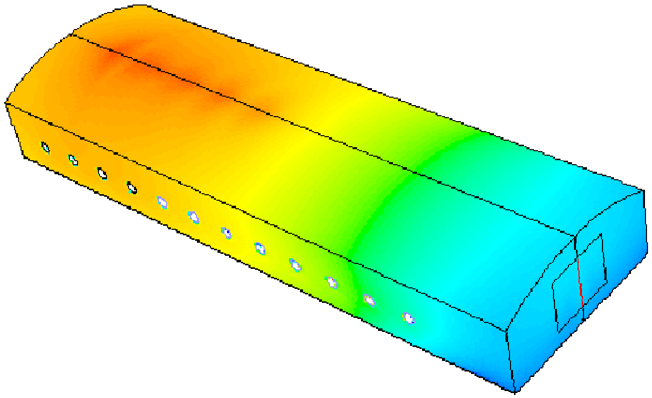

Casting with Water Cooling

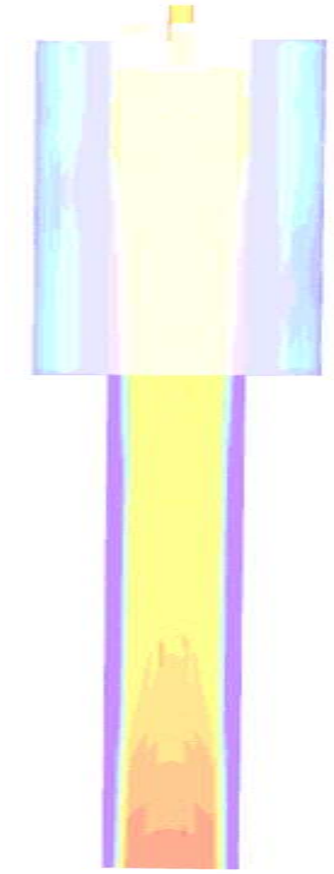

HVAC

NYC Urban Canyon Analysis

Ventilation Indoor Rain Forest

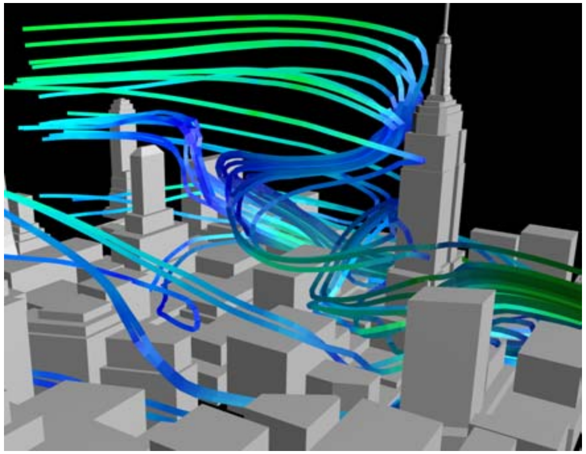

Dispersion Modeling

Cook Top

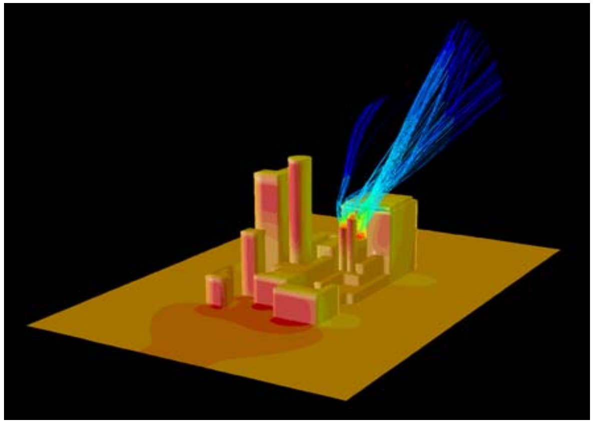

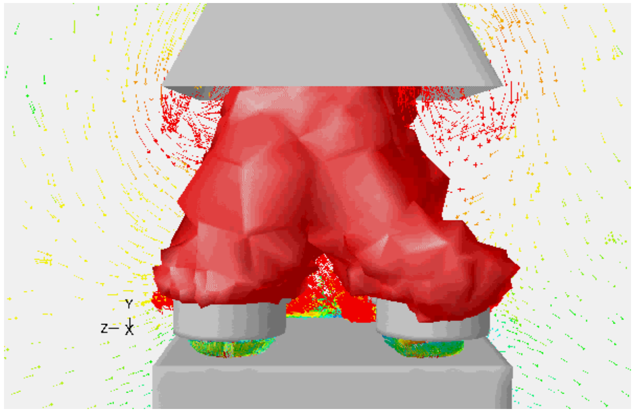

Consumer Products

Baby Care

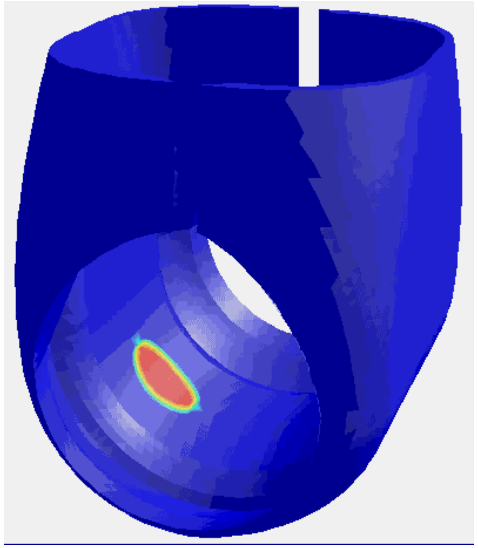

Diaper pad forming machine

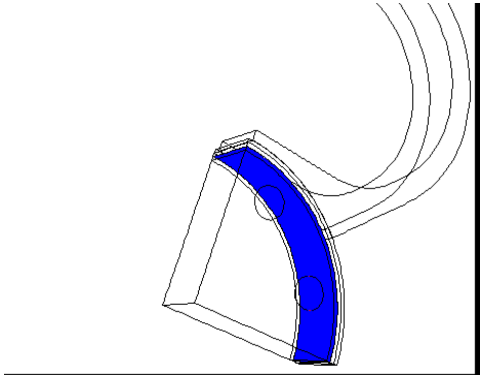

Food Processing

Pasta Extrusion

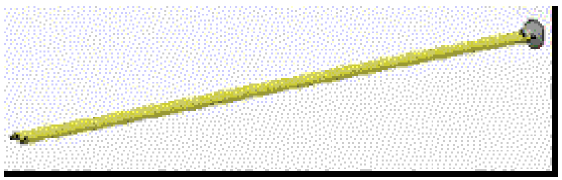

Dough Mixer

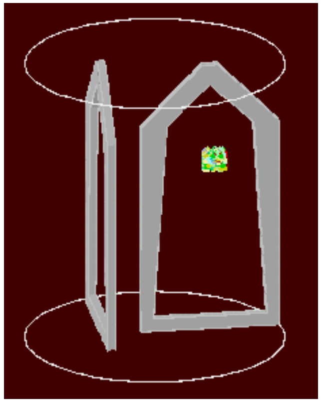

Brewing

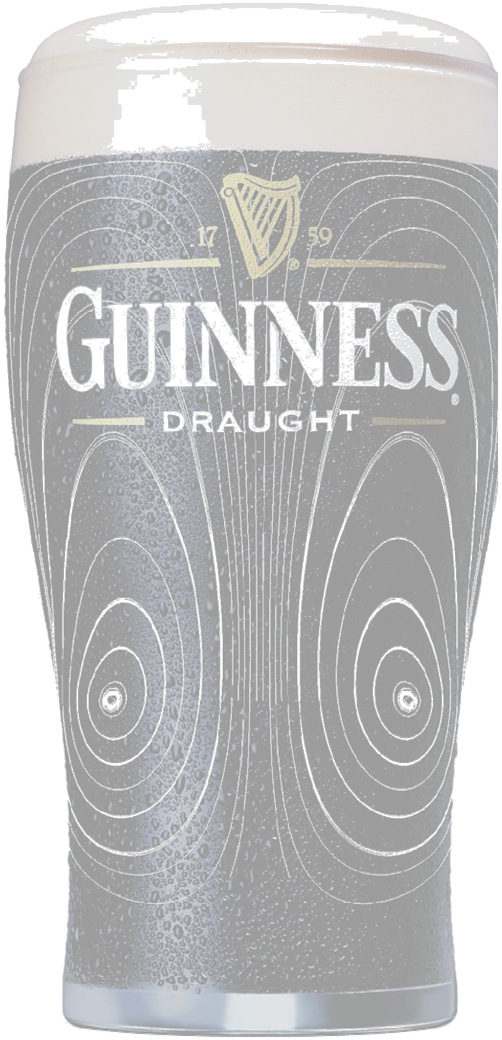

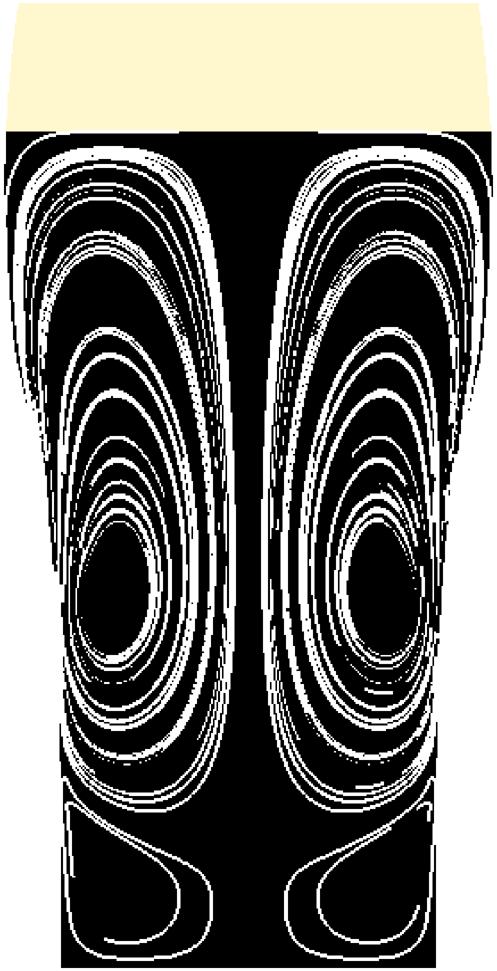

Major CAE Applications

- Implicit Structures
  - MSC.NASTRAN
  - ANSYS
  - ABAQUS
  - AMLS
- Explicit Structures
  - LS-DYNA
  - PAM-Crash
  - RADIOSS
- Computational Fluid Dynamics
  - Fluent, FIDAP
  - STAR-CD/HPC
  - PowerFLOW
  - CFX Family
  - Fire
  - Radioss-CFD
  - PAM-Flow

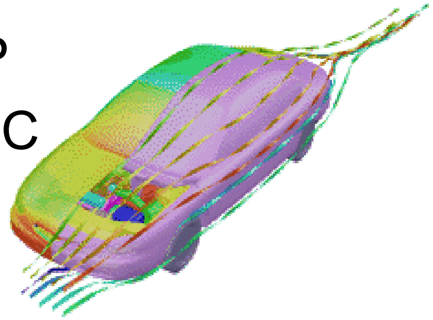

Factors that determine the proper hardware solution

- Application Availability
  - Architecture (IA32, IA64, etc.)
  - Operating System (HPUX or Linux)
  - Interconnects (version of MPI)
- Application Characteristics
  - SMP or Distributed Memory
  - Scaling
  - 32 or 64 bit
  - I/O requirements
- Problem Size
  - Memory Required (3 GB IA32)
  - File Size (2 GB IA32 Ext3)
  - Domain Decomposition

Structural Application Characteristics

Implicit Structures

- SMP systems
- In general these applications do not scale well (4 CPUs)
- High speed scratch area
- A 1 million DOF problem would require about 800 MB of memory
- A 5 million DOF problem would require about 2 GB of memory
- 100 GB scratch files

Application Availability

ISV	Application	IA64 HPUX	IA64 Linux	IA32 Linux
Abaqus	Abaqus Std	6.3-5	6.3-4	6.3
Adina R&D	Adina	8.0	8.0	8.0
Ansys	Ansys	7.0	7.0	7.0
CDH	AMLS	2.0	-	-

Application Sensitivity

Application	Memory Bandwidth	Cache size	I/O Bandwidth	Interconnect Latency	Interconnect Bandwidth
Abaqus Std	no	no	yes	NA SMP only	NA SMP only
Ansys	no	yes	yes	NA SMP only	NA SMP only
AMLS	no	no	yes	NA SMP only	NA SMP only
MSC.Nastran	yes/NVH no	no	yes	no	no
PAM-Crash	no	yes	no	yes	no
LS-Dyna	no	yes	no	yes	no
Radioss	no	yes	no	yes	no
STAR-CD	very	yes	some	yes	yes for big models
PowerFlow	slightly	yes	some	?	?
Fluent	yes	yes	some	yes	yes for big models

HP Systems for Implicit Structures

Server	Max CPUs	CPU Speed	Cache Size	Memory Bandwidth	Slot Type	Max Slots
rx2600	2	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	8.5 GB/sec	64 bit/133 MHz PCI-X	4
rx5670	4	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	12.8 GB/sec	3 @ 133 MHz PCI-X 6 @ 66 MHz PCI-X 1 @ 33 MHz PCI	10
Superdome	64	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	256 GB/sec	128 standard PCI-X slots 64 high BW PCI-X	192 with IOX
DL360	2	3.06 GHz	512 KB or 1 MB	4.2 GB/sec	PCI-X slots	2

High Speed Scratch

Fibre Channel

- Host Bus Adaptor
  - 2Gb FC Controller
  - A6795A
- Disk System 2405
  - 2 Gb Fibre Channel
  - 200 MB/sec
  - Dual Bus
  - 15 Disk Slots
  - 3U Rack Height
  - Rack Part number A6250AZ
- Disks
  - A6192A 36GB 10K rpm drive
  - A6193A 36GB 15K rpm drive
- Striping Software
  - Online JFS

MSC.NASTRAN Public Benchmarks XXCMD

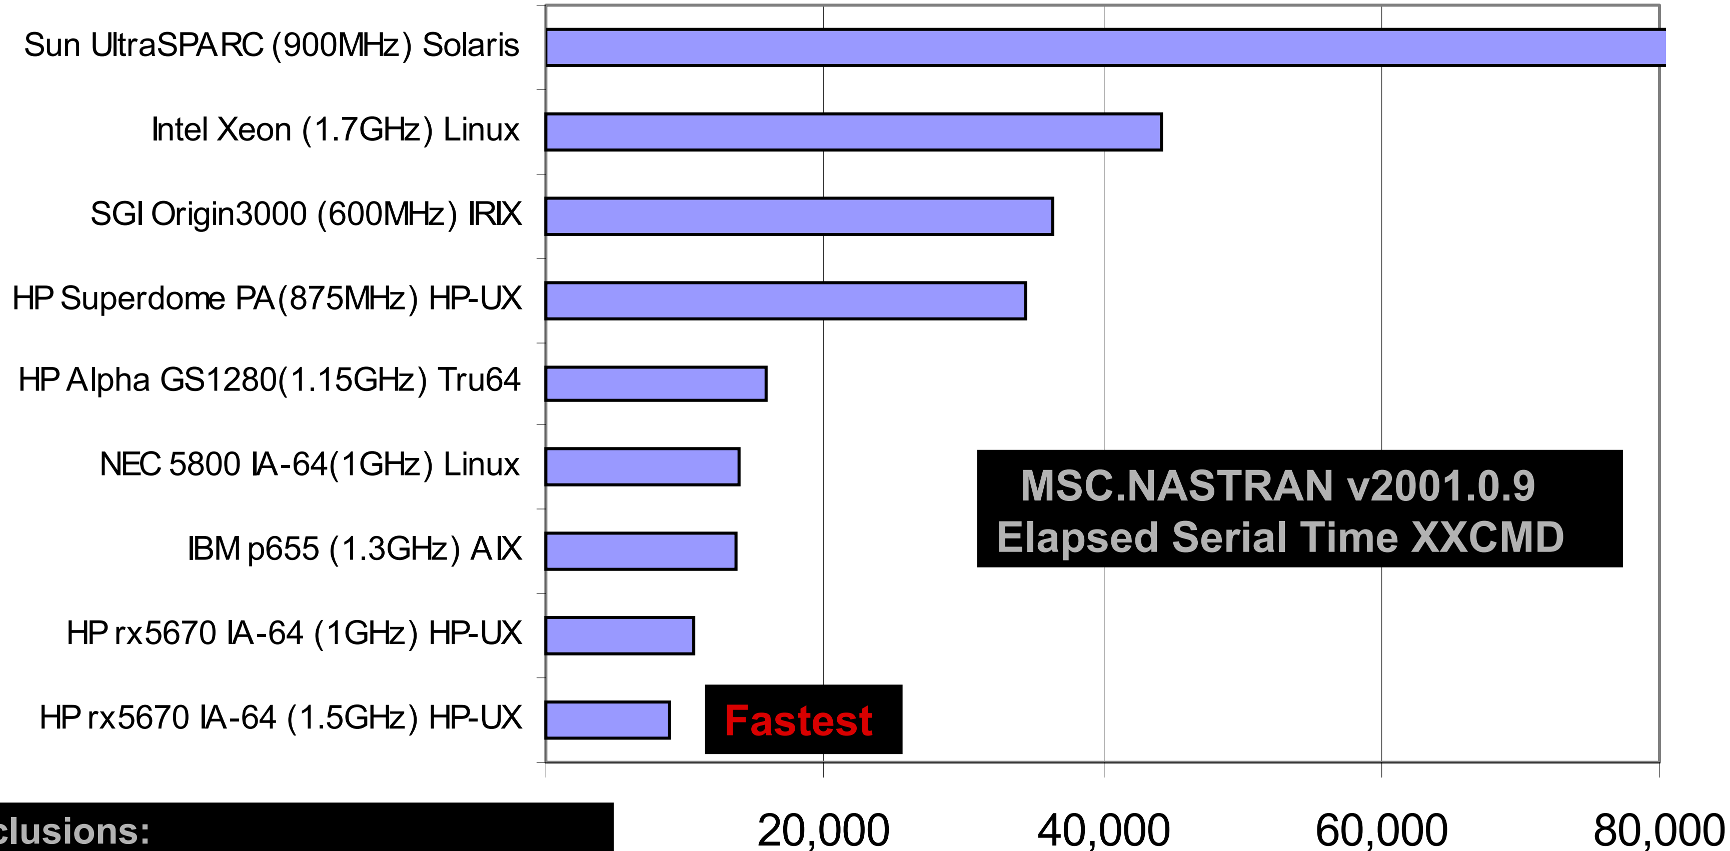

Conclusions:

- Itanium2 1GHz is 29% faster than IBM 1.3GHz
- HP Itanium2 running HP-UX is 31% faster than NEC Itanium2 running Linux, both at 1GHz

Reference: v2001 data posted at www.mscsoftware.com/support/prod_support/nastran/performance/index.cfm *as* of 3-Jul-03 except for 1.5GHz IA-64 HP-UX data. XXCMD is Solution 103 run on Car Body model with 1,584,000 DOF. HP realizes Intel IA-32 1.7GHz data plotted is old Xeon model but this is only IA-32 data posted by MSC. One of three vector architectures posted outperform Itanium2 but is in a different price class.

Structural Application Characteristics

HP WORLD 2003
Solutions and Technology Conference & Expo

Explicit Structures

- Moving from SMP to Clusters
- Scales well (ratio of computation to communication) depending on the amount of contact
- A 1 million DOF problem would require about 800 MB of memory
- A 5 million DOF problem would require about 2 GB of memory

Application Availability

ISV	Application	IA64 HPUX	IA64 Linux	IA32 Linux
ESI	PAM-Crash	v2002	-	v2002
LSTC	LS-Dyna	v970	v970	v970
Mecalog	Radioss	4.1q SPMD	4.1p SPMD	

Application Sensitivity

Application	Memory Bandwidth	Cache size	I/O Bandwidth	Interconnect Latency	Interconnect Bandwidth
Abaqus Std	no	no	yes	NA SMP only	NA SMP only
Abaqus Exp	no	no	no	NA SMP only	NA SMP only
Ansys	no	yes	yes	NA SMP only	NA SMP only
AMLS	no	no	yes	NA SMP only	NA SMP only
MSC.Nastran	yes/NVH no	no	yes	no	no
PAM-Crash	**no**	**yes**	**no**	**yes**	**no**
LS-Dyna	**no**	**yes**	**no**	**yes**	**no**
Radioss	**no**	**yes**	**no**	**yes**	**no**
STAR-CD	very	yes	some	yes	yes for big models
PowerFlow	slightly	yes	some	?	?
Fluent	yes	yes	some	yes	yes for big models

HP Systems for Explicit Structures

Server	Max CPUs	CPU Speed	Cache Size	Memory Bandwidth	Slot Type	Max Slots
rx2600	2	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	8.5 GB/sec	64 bit/133 MHz PCI-X	4
rx5670	4	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	12.8 GB/sec	3 @ 133 MHz PCI-X 6 @ 66 MHz PCI-X 1 @ 33 MHz PCI	10
Superdome	64	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	256 GB/sec	128 standard PCI-X slots 64 high BW PCI-X	192 with IOX
DL360	2	3.06 GHz	512 KB or 1 MB	4.2 GB/sec	PCI-X slots	2

Cluster Interconnects

- HPUX
  - Gigabit Ethernet
  - Hyperfabric
- Linux
  - Gigabit Ethernet
  - Myrinet
  - Quadrics

Radioss - Neon Benchmark

Mecalog

RADIOSS

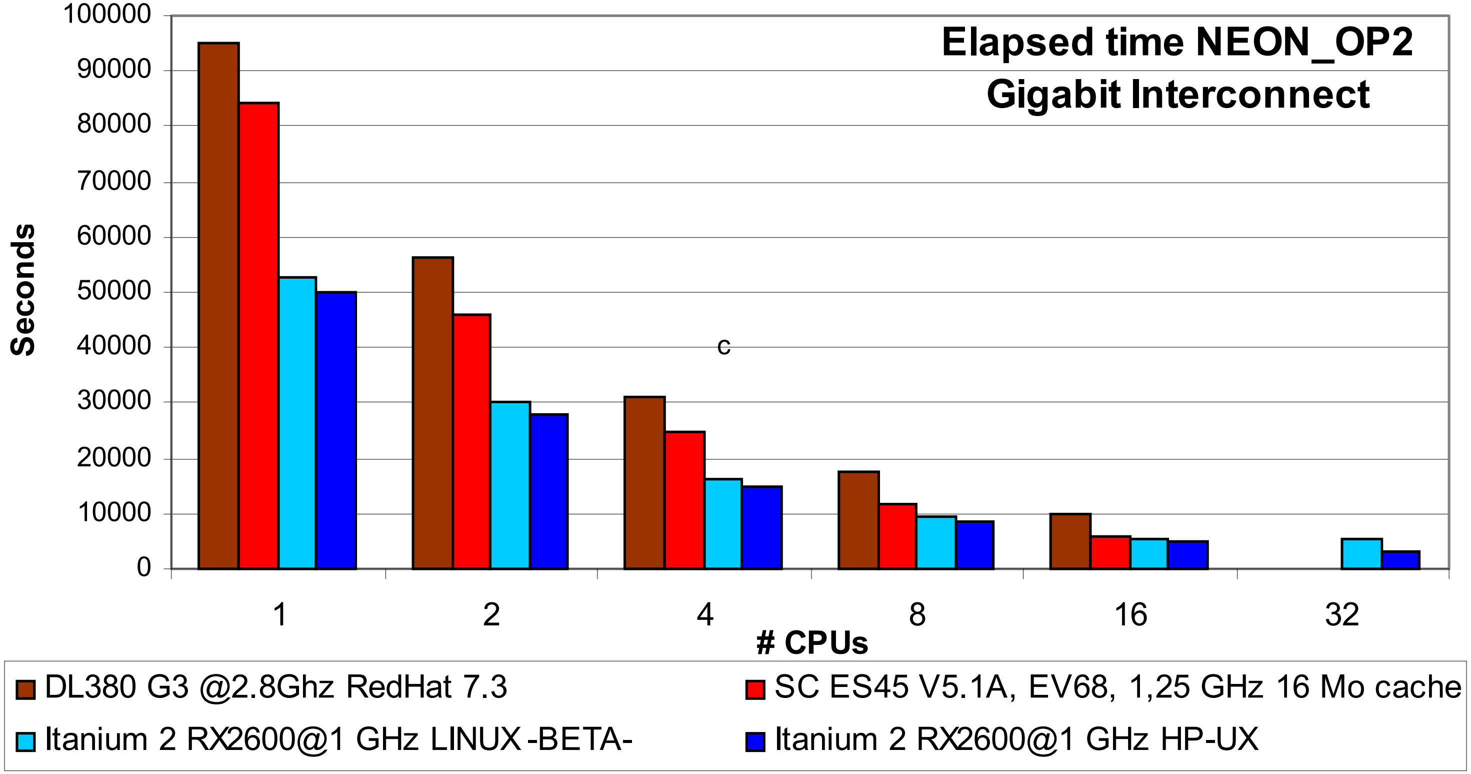

Conclusions:

- Itanium2 outperforms Alpha & Xeon nearly 2x on 1 CPU
- Itanium 2/HP-UX is ~10% faster than Itanium2/Linux
- HyperFabric shows improvement vs GigaBit at >32 CPUs

Reference: Neon_OP2 has 300,000 elements; run for 120ms

This data was run March'03 by Thierry Caron of HP on various HP systems.

CFD Application Characteristics

- Clusters are used more than SMP
- Scales well (ratio of computation to communication)
- Every million cells uses about 1 GB of memory
- 10 million cell problem in Fluent would use approximately 4 GB of memory using metis
- The amount of memory is larger when the problem is divided across multiple systems in a cluster

Application Availability

ISV	Application	IA64 HPUX	IA64 Linux	IA32 Linux
CD/Adapco	STAR-CD	3.14A	3.14	3.14
EXA	PowerFlow	3.4p2a		3.4p2a
Fluent	Fluent	6.1.22	6.1.22	6.1.22

Application Sensitivity

Application	Memory Bandwidth	Cache size	I/O Bandwidth	Interconnect Latency	Interconnect Bandwidth
Abaqus Std	no	no	yes	NA SMP only	NA SMP only
Abaqus Exp	no	no	no	NA SMP only	NA SMP only
Ansys	no	yes	yes	NA SMP only	NA SMP only
AMLS	no	no	yes	NA SMP only	NA SMP only
MSC.Nastran	yes/NVH no	no	yes	no	no
PAM-Crash	no	yes	no	yes	no
LS-Dyna	no	yes	no	yes	no
Radioss	no	yes	no	yes	no
STAR-CD	very	yes	some	yes	yes for big models
PowerFlow	slightly	yes	some	?	?
Fluent	yes	yes	some	yes	yes for big models

HP Systems for CFD

Server	Max CPUs	CPU Speed	Cache Size	Memory Bandwidth	Slot Type	Max Slots
rx2600	2	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	8.5 GB/sec	64 bit/133 MHz PCI-X	4
rx5670	4	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	12.8 GB/sec	3 @ 133 MHz PCI-X 6 @ 66 MHz PCI-X 1 @ 33 MHz PCI	10
Superdome	64	1.3 GHz 1.5 GHz	L1 32 KB L2 256 KB L3 3/6 MB	256 GB/sec	128 standard PCI-X slots 64 high BW PCI-X	192 with IOX
DL360	2	3.06 GHz	512 KB or 1 MB	4.2 GB/sec	PCI-X slots	2

Cluster Interconnects

- HPUX
  - Gigabit Ethernet
  - Hyperfabric
- Linux
  - Gigabit Ethernet
  - Myrinet
  - Quadrics

STAR-CD Public Benchmark

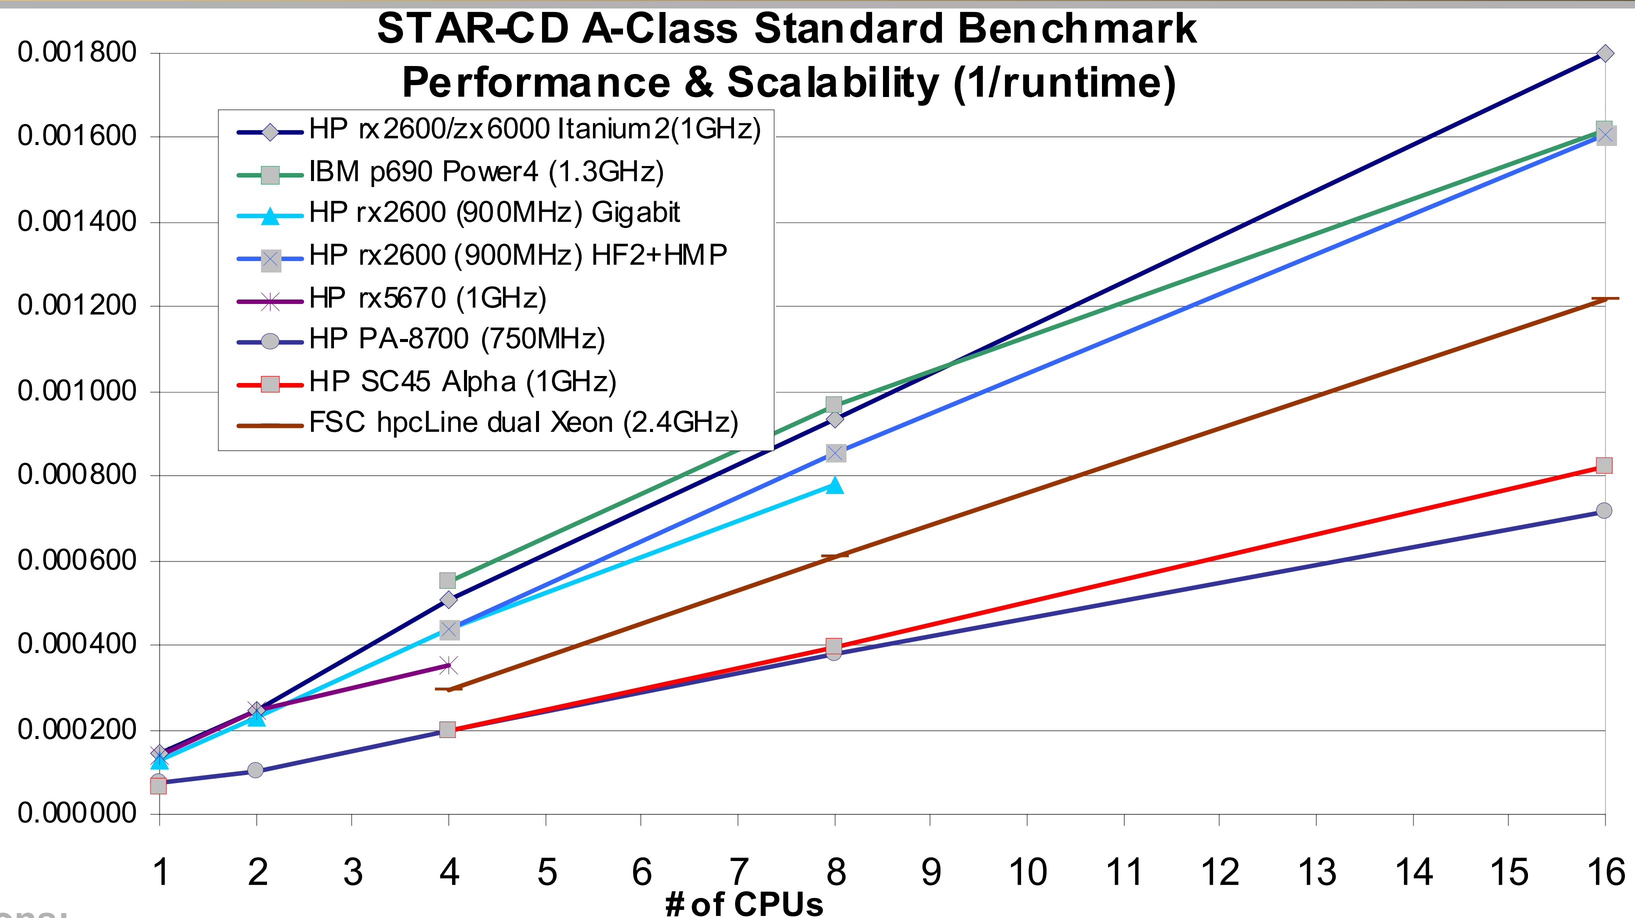

Conclusions:

- HyperFabric cluster of Itanium 2 (1 GHz) is 11% faster than Power4 (1.3GHz) SMP at 16 CPUs
- Lower cost Itanium2 (900MHz) is ~ equal at 16 CPU
- Itanium 2 is 48% faster than Xeon2.4GHz/Linux

Benchmark: Turbulant Flow around A-CLASS Car Dataset 5.9 Million Cells; SCALAR-CGS Solver. See: www.cd-adapco.com/support/bench/315/aclass.htm HP-UX Itanium2 data run with beta release STAR-CD.

Exa Public Benchmarks
optimized Itanium 2 results - April, 2003

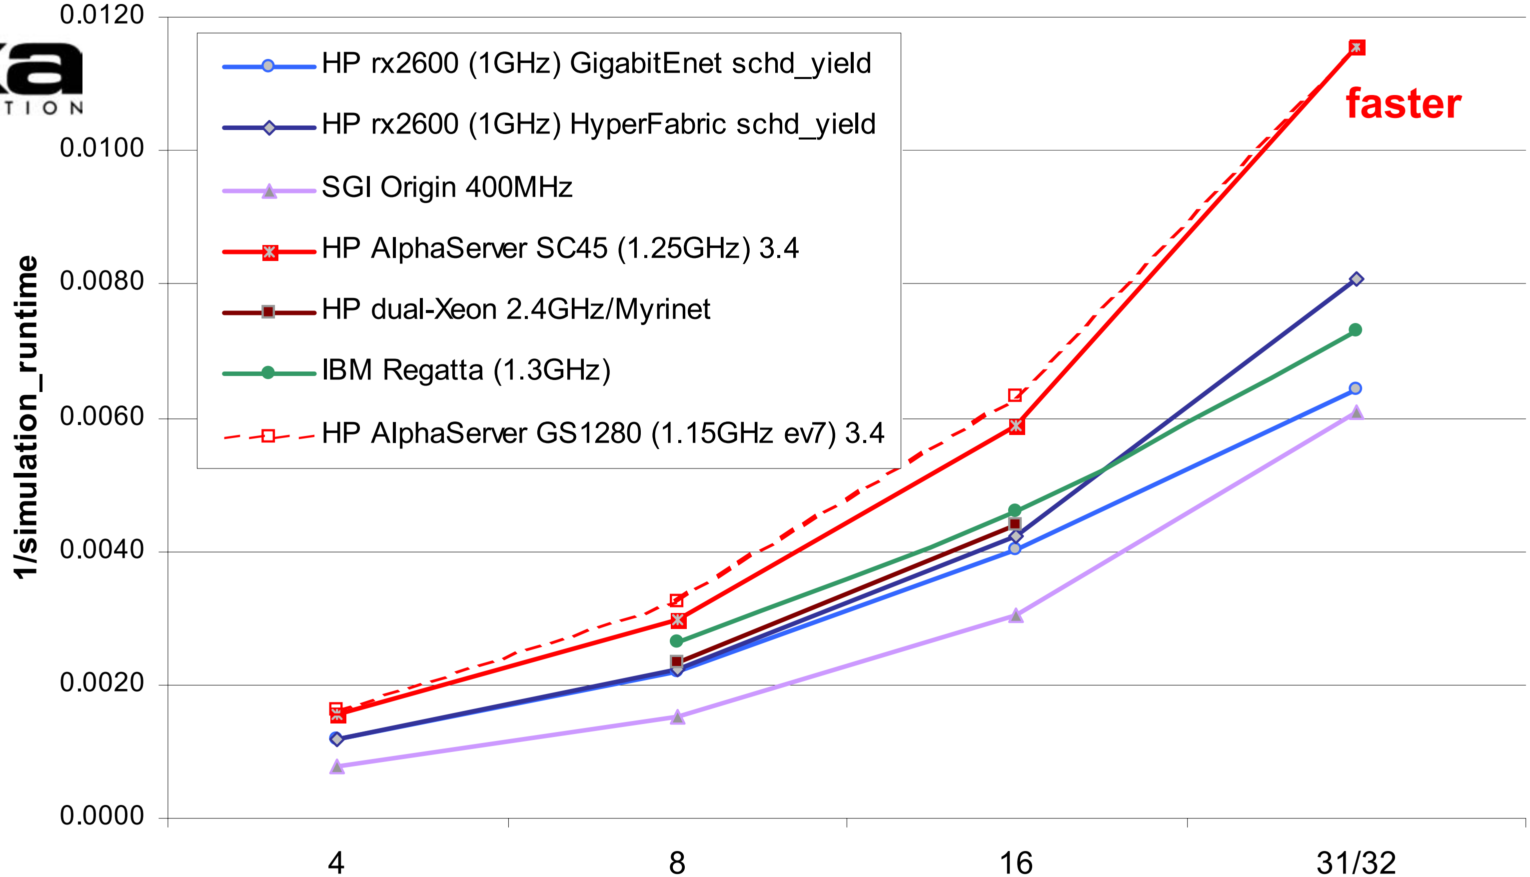

Observations:

- Itanium 2 is almost as fast as Linux 2.4 GHz at 16 cpus and likely outperforms it at 32 cpus
- HyperFabric2 gain over Gigabit is marginal at 8 & 16 cpu but 26% at 32-cpu due to lower latency

Reference: Exa benchmark posted at external password-protected ExaSource web site called external_case_1.
IBM 16-cpu datapoint is HP-interpolated data to curve-fit reported 8 & 32-cpu datapoints.

Conclusion

ONLY WITH AN UNDERSTANDING OF APPLICATION(S) CAN ONE TRULY DESIGN AN APPROPRIATE SOLUTION

Resources

How to reach the Americas team

Virtual LIFE SCIENCES TEAM Contacts

Doug Brown, Industry Segment Lead

512 418 1595 douglas.brown@hp.com

Jon Saposhnik, Business Development

949 498 3406 jon.saposhnik@hp.com

Joel Jaffe, Business Development

603 884 2367 joel.jaffe@hp.com

How to reach the Americas team

Virtual CAE TEAM Contacts

Dave McDonnell Industry Segment Lead

905 206 3108 dmcdonnell@hp.com

Dan Fennell, Business Development

612 326 2063 dan.fennell@hp.com

Greg Worth, Business Development

248 651 1673 greg.worth@hp.com

Customer Web Resources for Life Sciences

Solutions for Life Sciences	http://www.hp.com/techservers/life_sciences/index.html
Applications	http://www.hp.com/techservers/life_sciences/applications.html
Benchmarks	http://www.hp.com/techservers/life_sciences/benchmarks.html
White Paper	http://www.hp.com/techservers/life_sciences/overview.html
HP & Oracle in Life Sciences	http://www.hp.com/techservers/life_sciences/sb_oracle.html
Technical Solutions	http://www.hp.com/technicalsolutions/life_sciences/

Customer Web Resources for CAE

Solutions for CAE	http://www.hp.com/techservers/cae/index.html
Applications	http://www.hp.com/techservers/cae/applications.html
Benchmarks	http://www.hp.com/techservers/cae/benchmarks.html
White Paper	http://esp.sell.hp.com:2000/nav24/ppos/os/itanium/TechWP/090017ad81605da8/
HP CAE News	http://www.hp.com/techservers/cae/index.html#news
CAE Events	http://www.hp.com/technicalsolutions/cae/index.html#events

Internal Web Resources for Life Sciences

HPTC Sales	http://esp.mayfield.hp.com:2000/nav24/ext/hptc/sr/index.htm
HPTC Americas	http://americashptc.americas.cpqcorp.net/mktdev/
Presentation Materials	http://esp.mayfield.hp.com:2000/nav24/ext/hptc/sr/solutions/index.htm#ls
Application Availability	http://techmktg.rsn.hp.com/davem/projects.htm
Benchmarks	http://techmktg.rsn.hp.com/davem/Benchmark_Data.htm
Market Intelligence	http://marketintelligence.corp.hp.com/homePres.asp
Life Sciences Advertising	http://esp.mayfield.hp.com:2000/nav24/ppos/servers/gen/PrAdsImages/LSMayAD.pdf

Internal Web Resources for Cluster Configuration Tools

HPTC Americas Configuration Tools	http://Americashptc.americas.cpqcorp.net/LINUX-clusters/LINUX-clusterblocks.htm
LC Tools and Configuration Guide	http://H18004.www1.hp.com/solutions/enterprise/highavailability/linux/highperformance/index.html
HPTC Americas	http://americashptc.americas.cpqcorp.net/mktdev/

Q&A

HP WORLD 2003
Solutions and Technology Conference & Expo

HP WORLD 2003

Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

hp

invent

Backup Slides

Resources

- Computer-aided engineering (CAE) applications and their suitability on Itanium-based platforms from HP http://esp.sell.hp.com:2000/nav24/ppos/os/itanium/TechWP/090017ad81605da8/
- ClusterBlocks
- Application List
- Yih Yih's paper
- LC site
- XC site

Performance Backup Slides

Itanium2 1.5GHz relative to 1.0GHz in systems

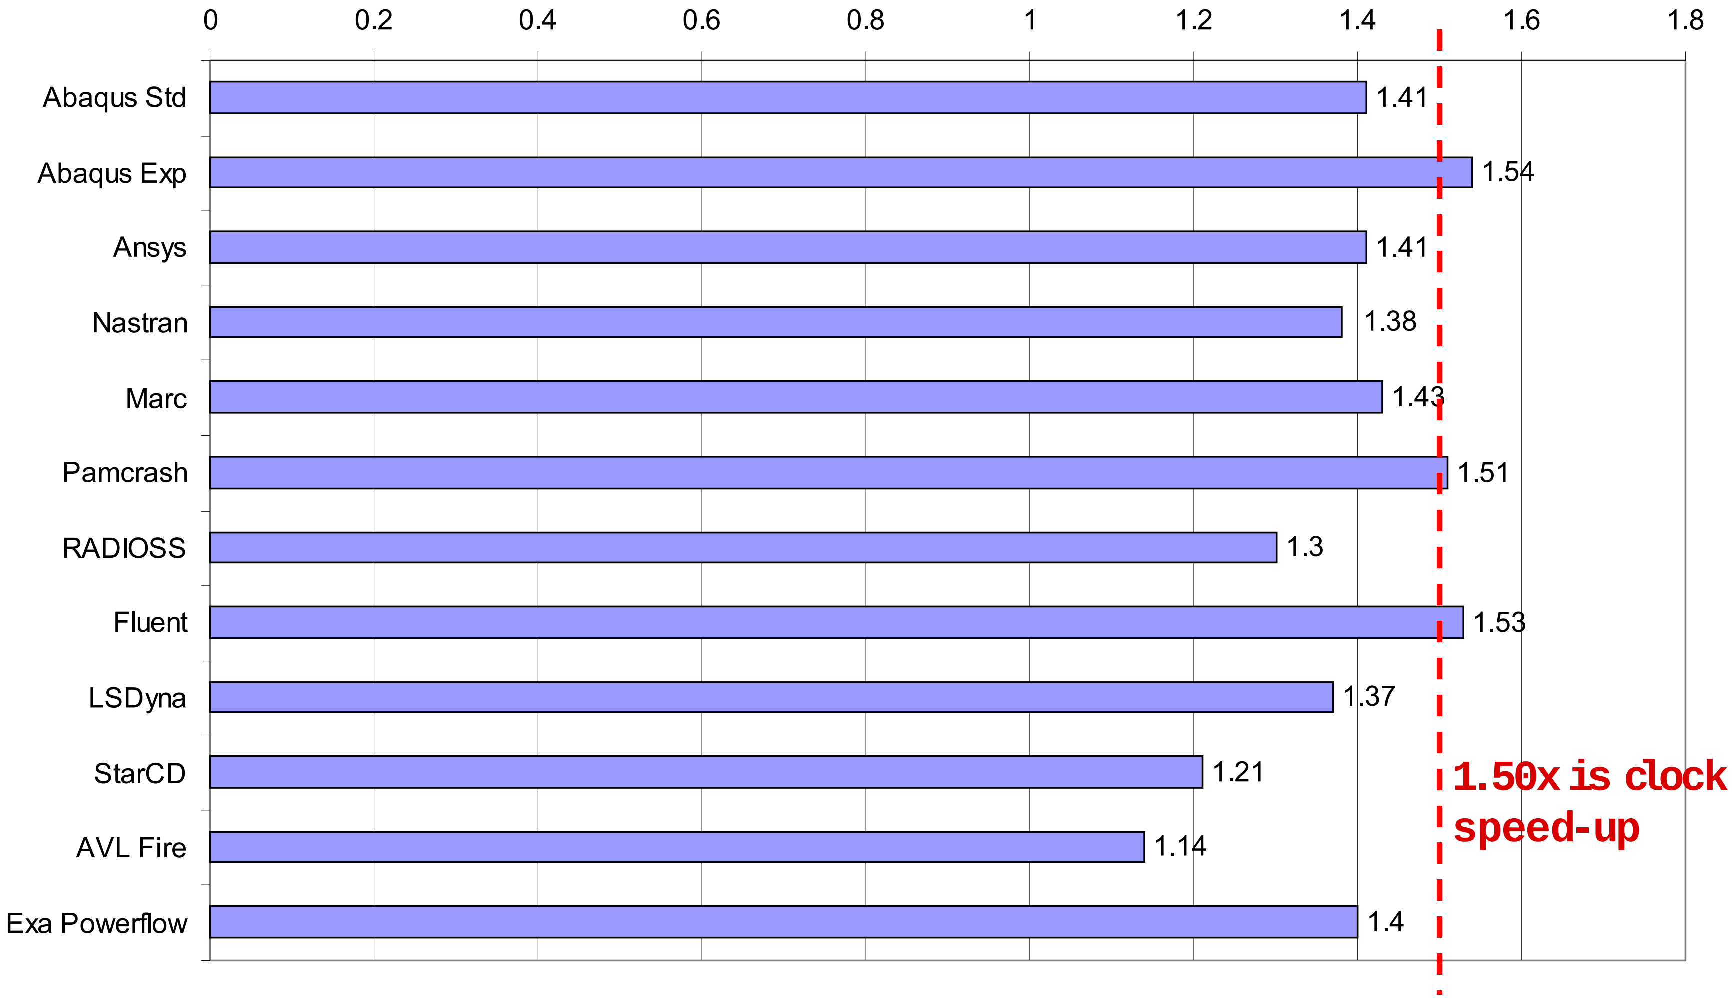

ISV applications run 1.4X to 1.5X faster on 1.5GHz/6MB-cache CPU's than with 1GHz/3MB-cache CPU's UNLESS the application is limited by memory bandwidth, in which case the performance ratio is ~1.2X

LS-DYNA NCAC Benchmark

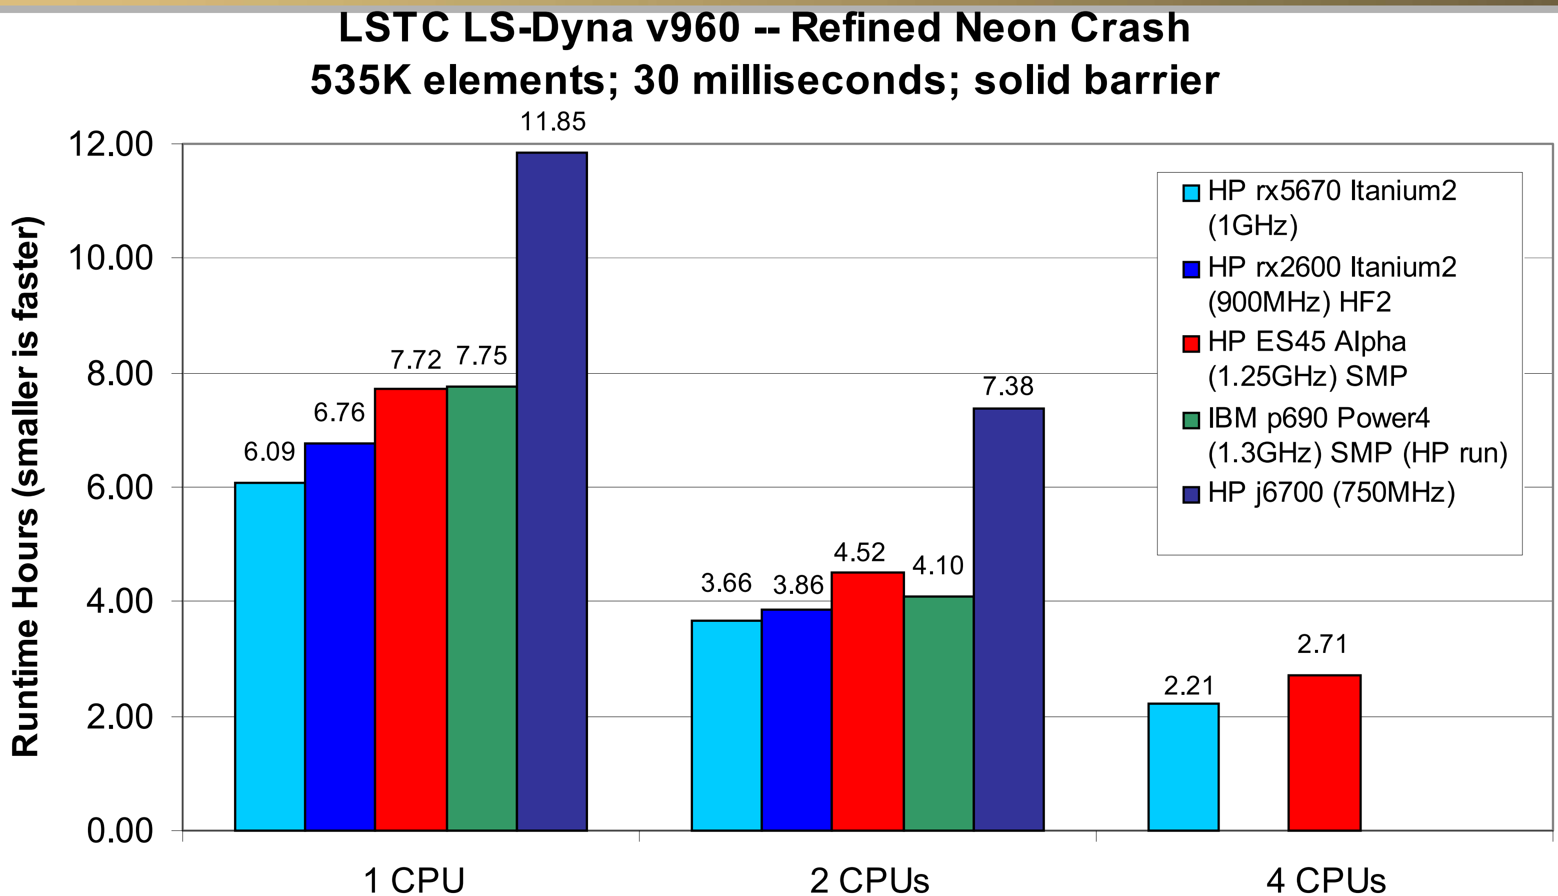

Conclusions:

- Itanium2 outperforms Alpha & Power4 for SMP Dyna
- Itanium 2 is 91% faster (1.91x) than PA-8700 at 2 CPUs
- speed-up from 900MHz to 1GHz is linear (11%) and near-linear speedup is expected when Madison is introduced

Reference: National Crash Analysis Center offers a model of Neon car for public benchmarking use: www.ncac.gwu.edu/research/Modeling/neon/Neon.html
This data was run by HP on symetric multiprocessor servers. No Intel/Xeon data available.

PAM-CRASH Performance

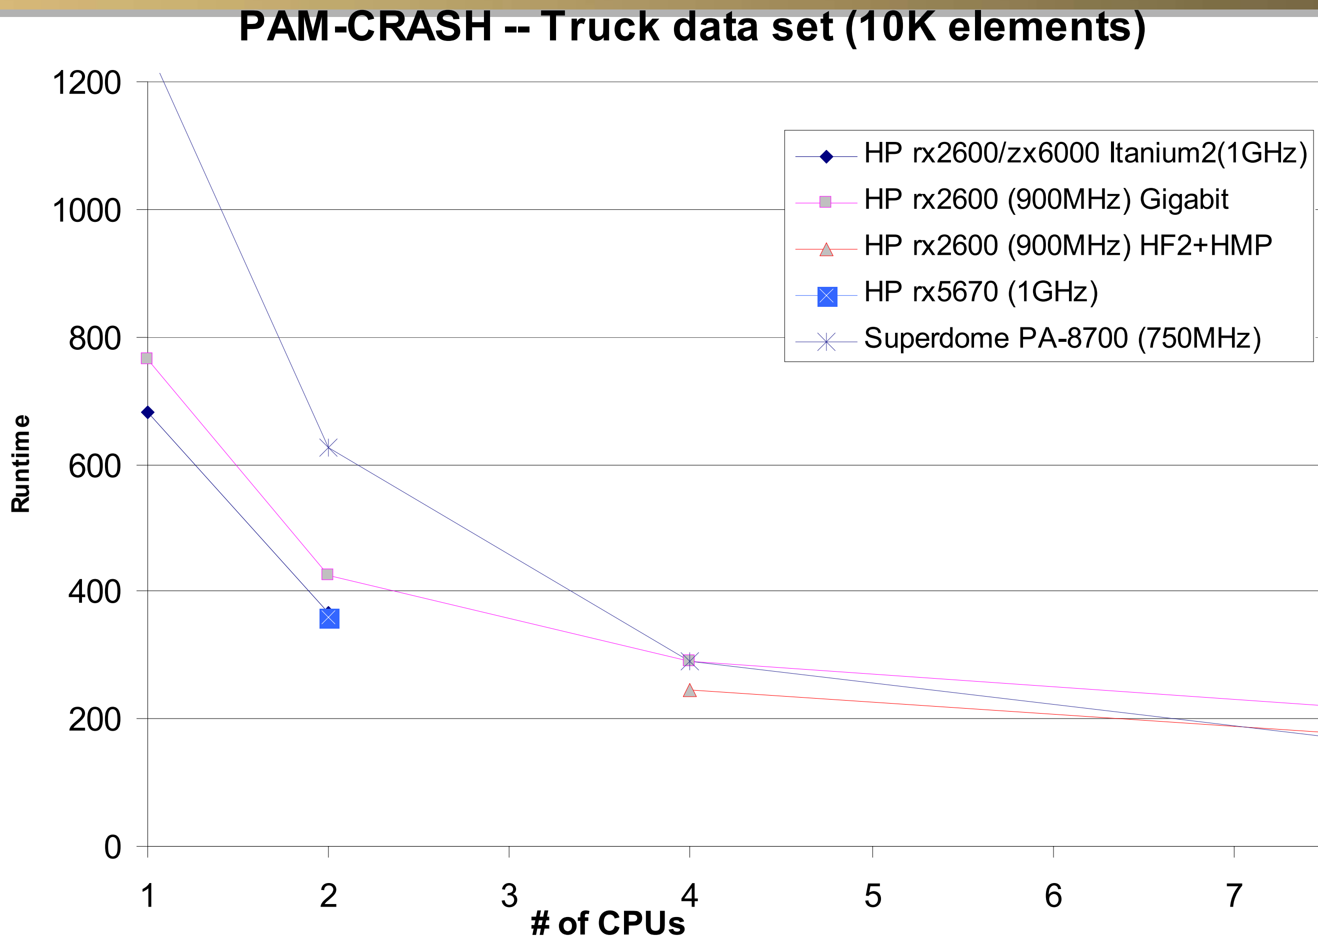

Conclusions:

- Itanium2 1GHz is 16% faster than 900MHz at 2CPUs, yet 900MHz can be much less $
- clustering capability of PAM-CRASH enables large cost effective systems to be constructed with Itanium 2 nodes

Reference: v2002 and HP-UX 11 used for all runs.
Data set is frontal crash, 10K elements, provided by ESI Group.
HP-UX/Itanium2 results submitted/approved by ESI Group.
For more information: www.hp.com/go/esi www.esi-group.com

ANSYS Public Benchmarks

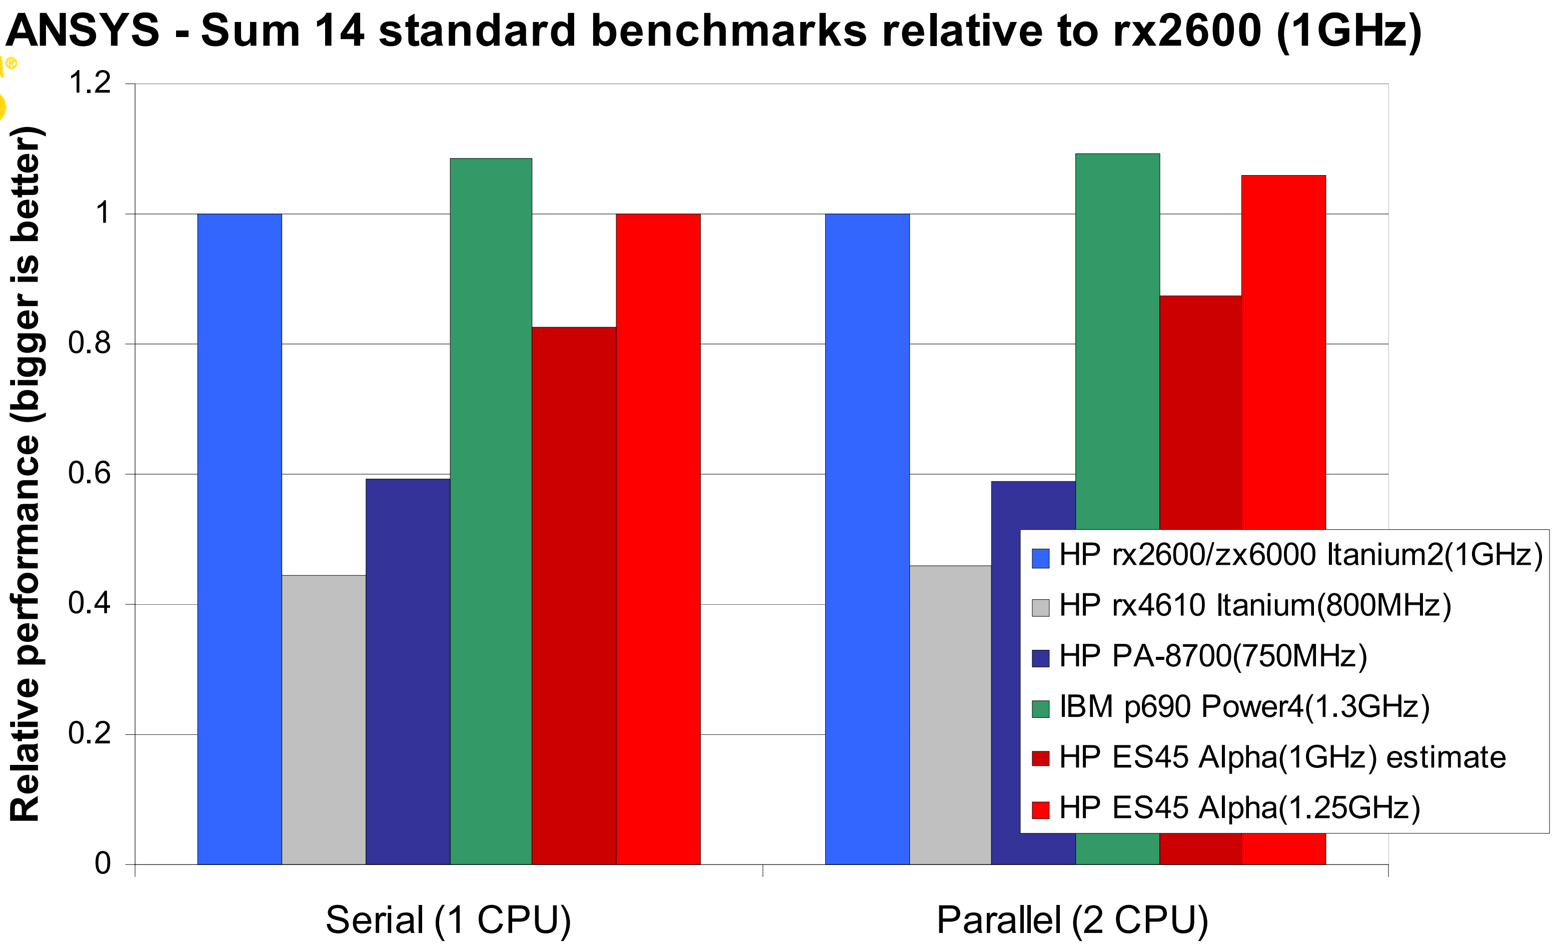

Conclusions:

- Itanium 2 averages 1.68x faster than PA-8700
- Itanium 2 almost as fast as Alpha and Power4 but at lower price

Reference: ANSYS 6.1 data posted at *www.ansys.com/services/hardware_support/61index.htm*
HP-UX Itanium2 data run with pre-release ANSYS 7.0 and results submitted/approved by ANSYS. All HP runs with HP-UX 11.x.

MSC.NASTRAN (sol 108) LGQDF

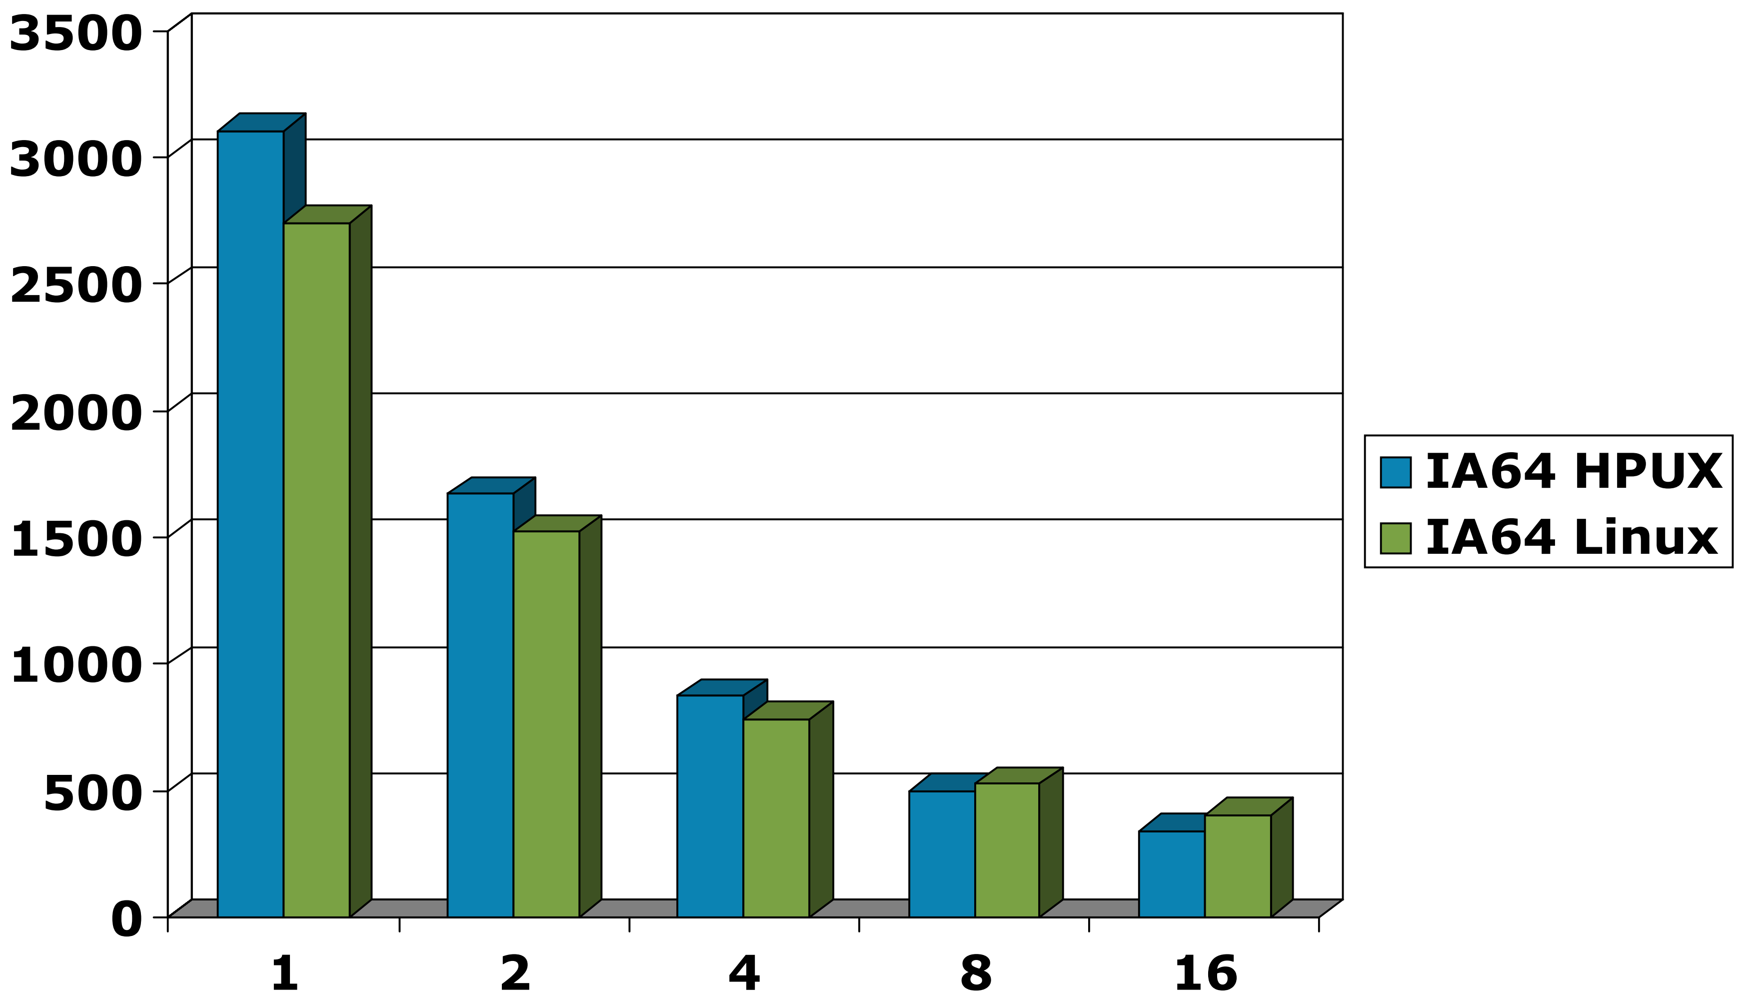

MSC.NASTRAN (sol 108) XLTDF

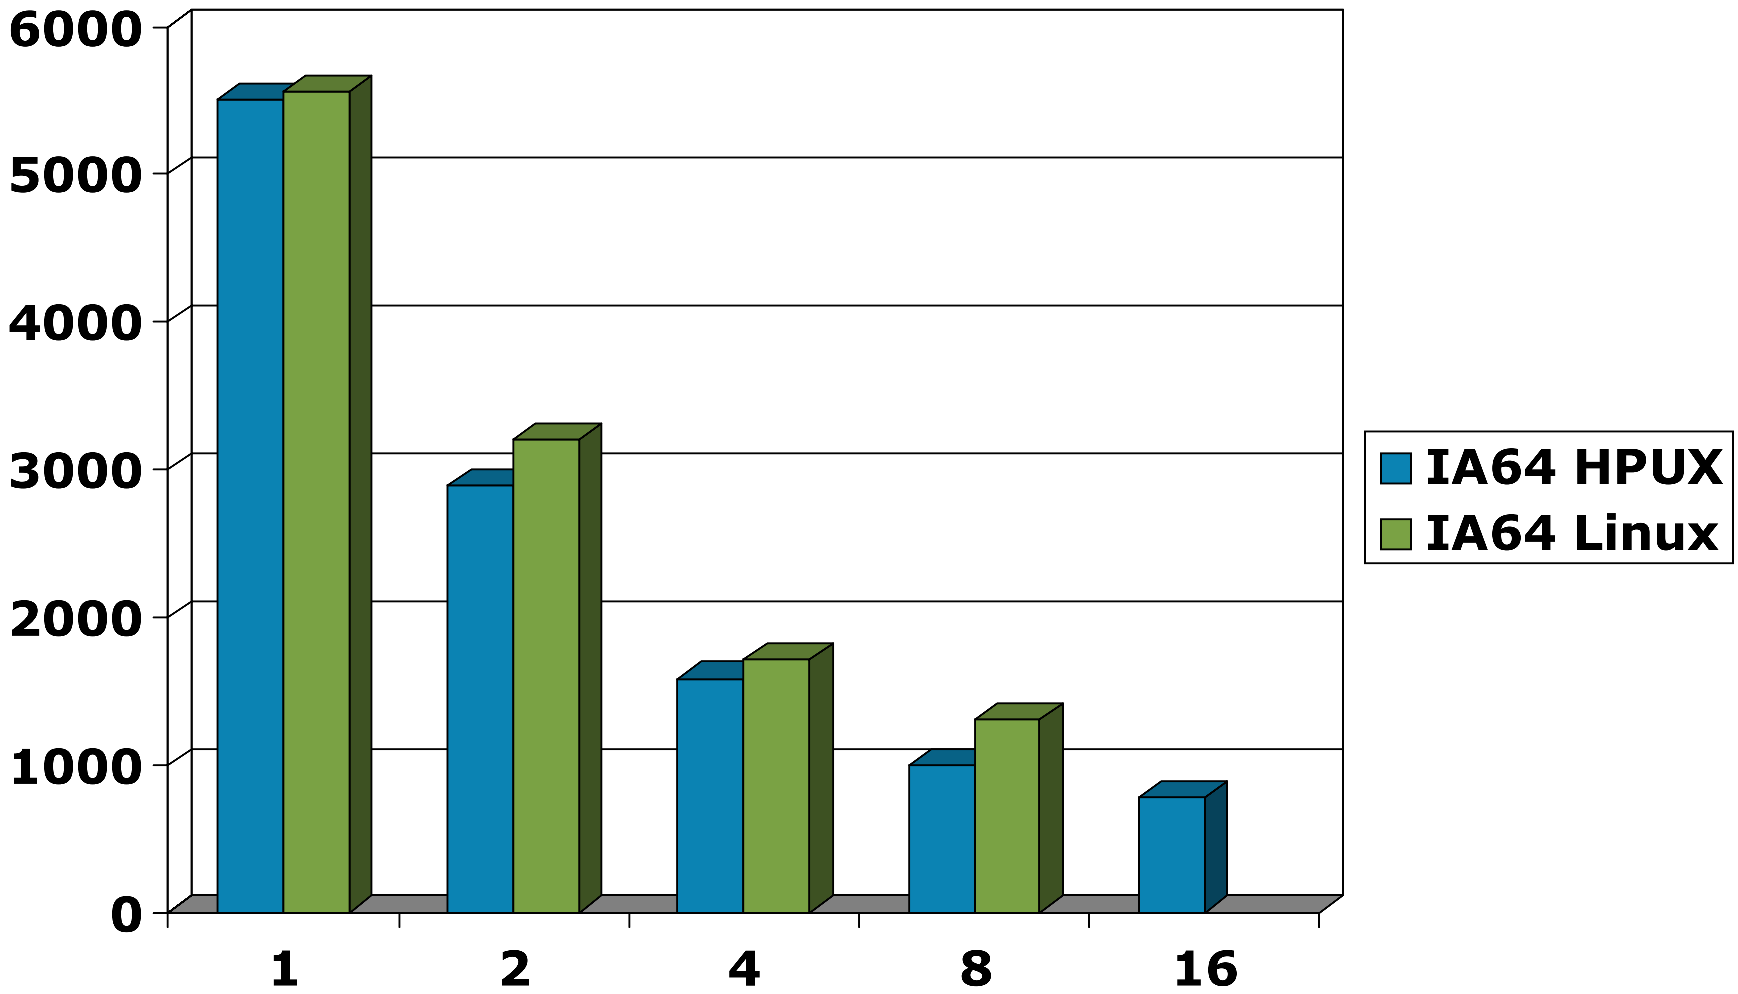

MSC.NASTRAN (sol 103) XXCMD

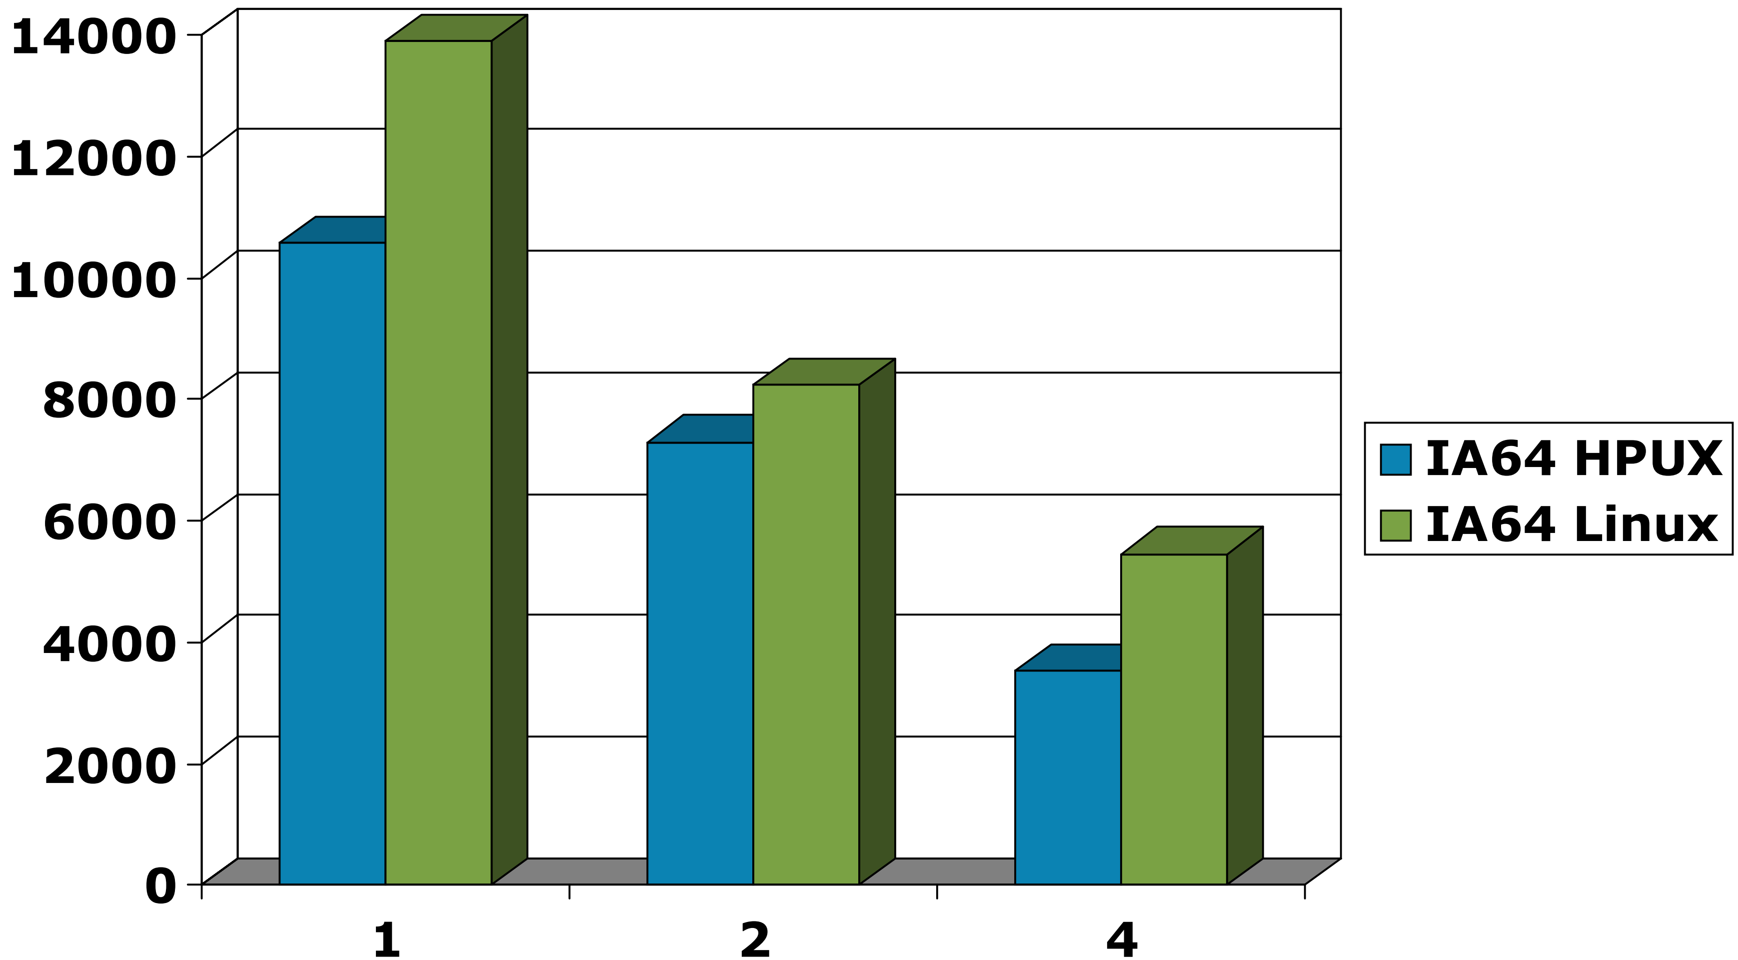

PAM-FLOW Public Benchmark

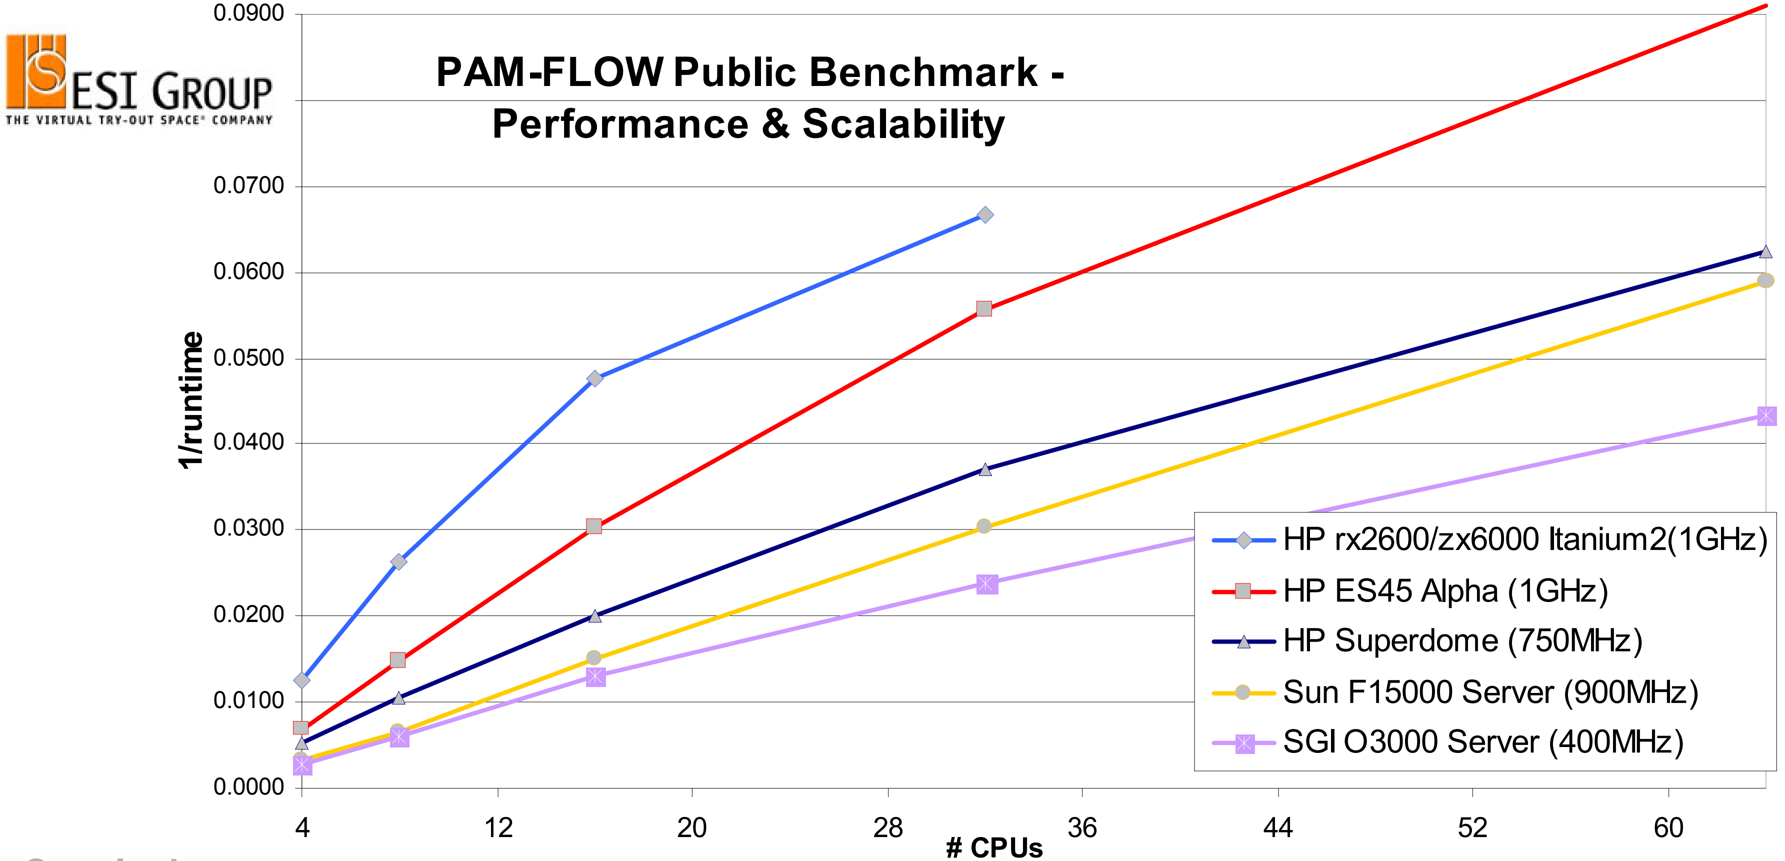

Conclusions:

- Itanium2 cluster is fastest system posted on ESI web site for PAM-FLOW benchmark at 32 CPUs (not tested by HP at 64 CPUs yet)
- HP Alpha system is also very fast and scalable

Reference: Data set is external aerodynamic public benchmark (mesh size of 3.7 million tetrahedral grids) posted by ESI Group at *www.esi-group.com/products/pamflow/index.php?page=perf*
On 1/6/03, no data posted for IBM Power4, 600MHz SGI, or 1.25GHz Alpha. For more information: www.hp.com/go/esi

FIRE 7.3 Performance

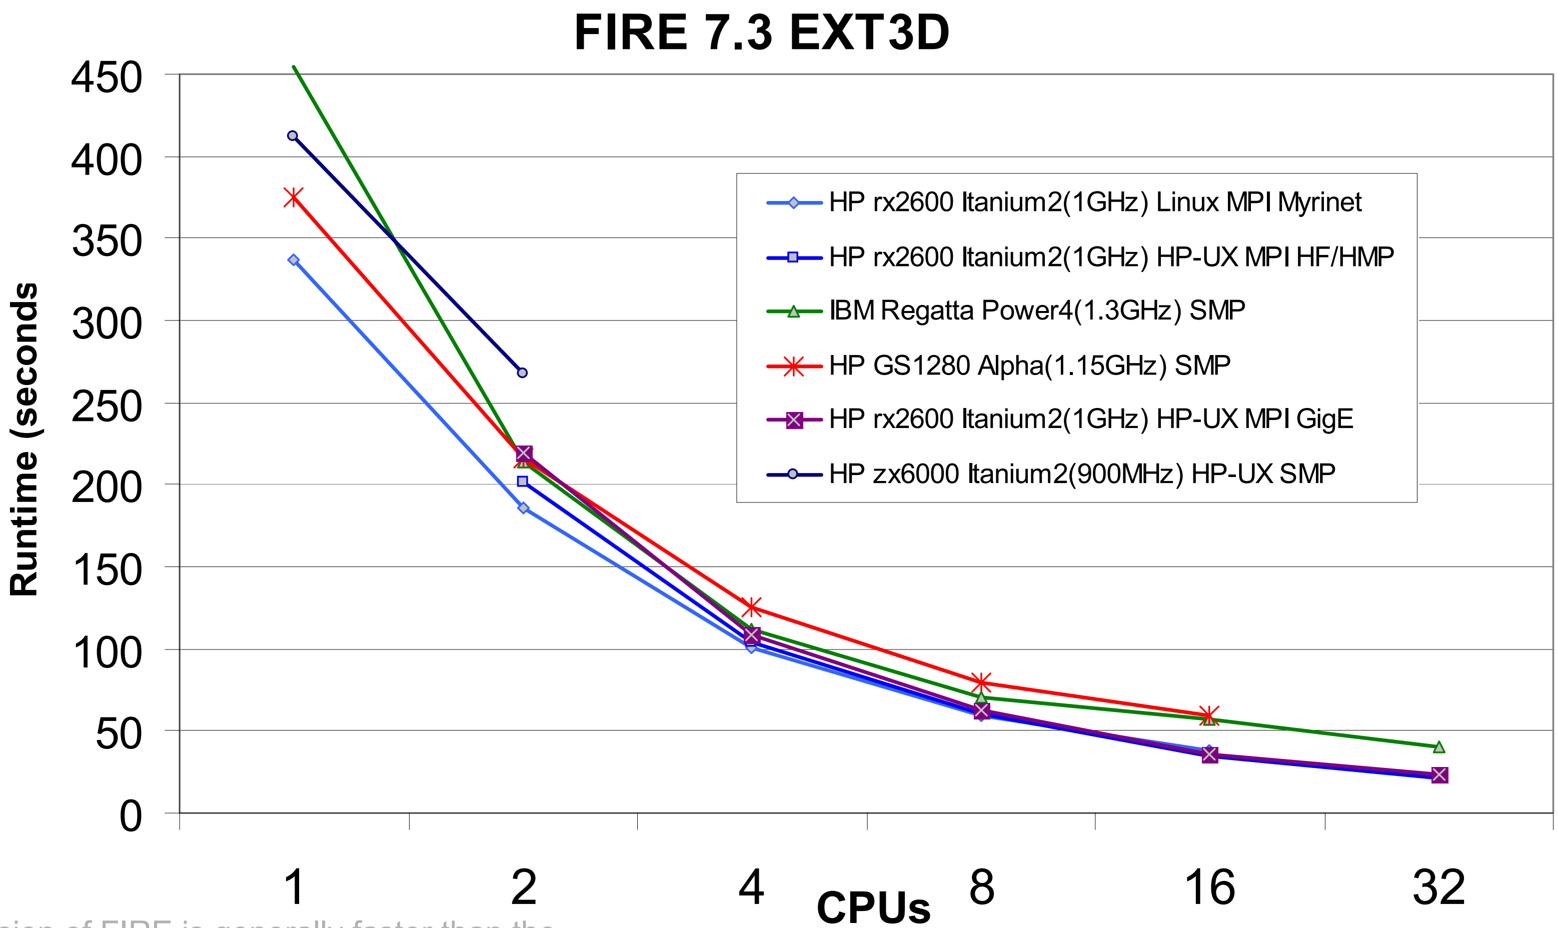

Notes:

- the MPI version of FIRE is generally faster than the SMP version, and new release is fully functional
- the benefit of a low latency interconnect begins to be apparent at 8 to 16 CPUs
- HP optimization efforts with HP-UX and F90 continue

Benchmark: EXT3D Car Aerodynamics 711K cells; unstructured mesh; 50 Iterations; TVD Differencing. Test Report available from: guenter.bachler@avl.com All data plotted are from April 2003 report.

FLUENT 1-CPU Relative Performance

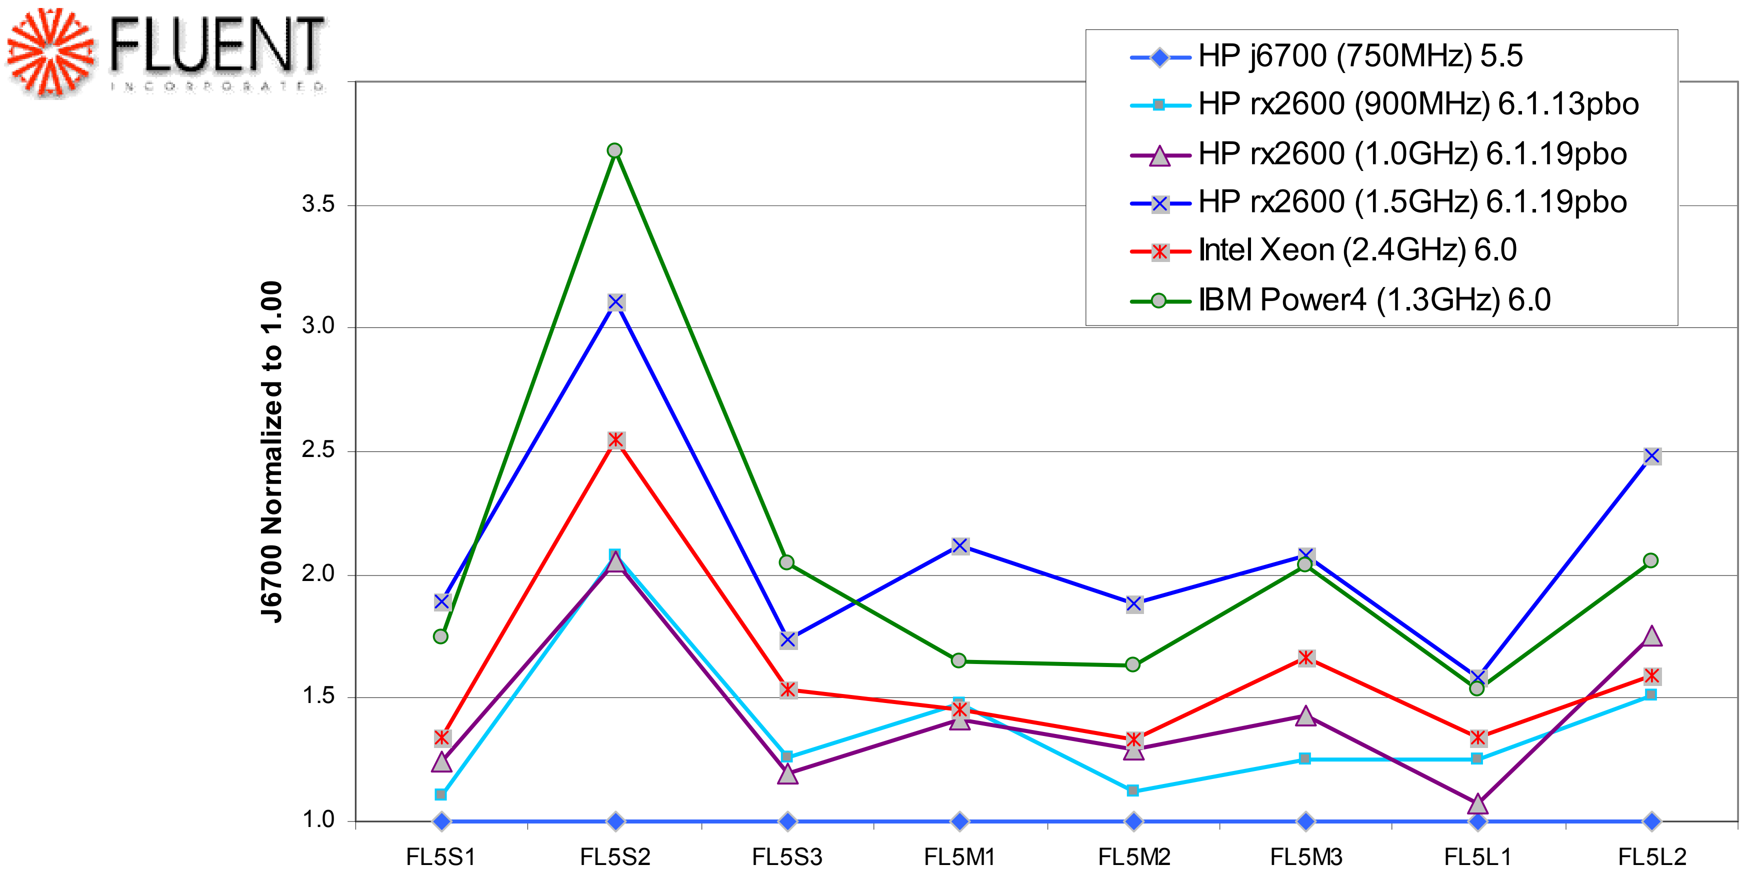

Observations:

- Itanium2 cluster for FLUENT offers 1.2 to 1.5x gain over PA-RISC and full 64-bit address space for grid generation >2 mil cells

Reference: Data set is transonic flow around a fighter aircraft, 848K cells, public benchmark. Data posted by Fluent at: www.fluent.com/software/fluent/fl5bench/index.htm except Itanium2 data. Fluent 6.1.18 (Feb03) fully supports Itanium2/HP-UX but production release in 6/03 will offer this optimized performance level.

Storage Backup Slides

High Speed Scratch (Ultra SCSI)

- Host Bus Adaptor
- Disk System 2110
  - Single Bus
  - U320 data transfer rates (320 MB/s)
  - 4 Disk Slots
  - 1U Rack Height
  - Rack Part number A7381AZ
- Disks
- Striping Software
  - Online JFS
  - LVM

High Speed Scratch (Ultra SCSI)

- Host Bus Adaptor
  - U160
  - Data transfer rate of 160 MB/s
  - Channel
  - A6829A/A7060A
- Disk System 2100
  - U160
  - Single Bus
  - Data transfer rate of 160 MB/s
  - 4 Disk Slots
  - 1U Rack Height
  - Rack Part number A5675AZ
- Disks
  - 36.4 GB 10K or 15K rpm or 73.4 GB 10K rpm drives
  - A6538A 36GB 10K rpm drives
  - A6541A 36GB 15K rpm drives
- Striping Software

High Speed Scratch (Ultra SCSI)

- Host Bus Adaptor
  - U160
  - A6829A/A7060A
- Disk System 2300
  - U160 enclosure upgradeable to U320
  - Single or split Bus
  - Data transfer rate of 160 MB/s
  - 14 Disk Slots
  - 3U Rack Height
  - Rack Part number A6490AZ
- Disks
  - A6538A 36GB 10K rpm drive
  - A6541A 36GB 15K rpm drive
- Striping Software

Stripe params

```
vgcreate –s 64 /dev/vg01 c1d1
vgextend /dev/vg01 c2d1 c3d1 c4d1 c1d2 c2d2 c3d2 c4d2 c1d3 c2d3
                    c3d3 c4d3 ….
lvcreate –i #nd –I 64 –n lvol1 /dev/vg01
```

Linux Backup Slides

Cluster Features

Feature	Requirement	Want	Not Needed
Head node mirrored OS			
Redundant head node			
I/O Node			
Redundant I/O paths			
Login node			
Job queuing			
High speed I/O			
Remote node power on/off			
Update management			
Power monitoring			
Temperature monitoring			
Outside application server			

Linux Issues

- Versioning
- Cluster Management
- Maximum memory
- Maximum file size

General info Linux

- Units:
  - 2^{10} bytes = 1 KiloByte
  - 2^{20} bytes = 1 MegaByte
  - 2^{30} bytes = 1 GigaByte
  - 2^{40} bytes = 1 TeraByte
  - 2^{50} bytes = 1 PetaByte
  - 2^{60} bytes = 1 ExaByte

- size of C types using gcc:

arch (sizeof)	char	int	long	longlong	double	float
ia64	1	4	8	8	8	4
ia32	1	4	4	8	8	4

Linux

Architecture	Memory Addressible by Linux Kernel	User Memory per Process	Comments
IA32	2^{32} = 4 GB 2^{36} = 64 GB	3 GB	(Without PAE36 enabled kernel, but PAE==>3-5% performance hit penalty) (But, going farther than 16 Gbytes may not work, PTE is too big for the first GigaByte) (per process maximum contiguous 2.1 GB, then 0.7GB in several sized chunks)
IA64	2^{41} = 2 TB 2^{46} = 64 TB		

Linux File System

	Max File Size	File System Limits Ext2/3	File System Limits xfs	NFS	
IA32	2^{32} = 2 GB 2^{63} = 8 EB	2 TB blocksize =4k	8 EB	2 GB V2 2^{64} V3	
IA64	2^{63} = 8 EB	64 TB blocksize =8k	8 EB		

Linux File Systems

	Partition Limit	VFS Limit	LVM Limit		
IA32	1 TB	16 TB	1 EB using 16Gbytes Physical Extent		
IA64	1 TB	16 TB	1 EB using 16Gbytes Physical Extent		

Linux File Systems

	Minix FS	Ext FS	Ext2 FS	Xia FS
Max File System	64 MB	2 GB	4 TB	2 GB
Max file Size	64 MB	2 GB	2 GB	64 MB
Max File Name	16/30 characters	255	255	248
Extensible	no	no	yes	no
Variable Block Size	no	No	yes	no
Maintained	yes	no	Yes	?

- With the advent of the Pentium Pro, it became apparent that support for more than 4GB of physical memory would be necessary. This led to the *Physical Address Extension* scheme developed by Intel. As a result, the PPro (and the subsequent P6 processors Pentium II and Pentium III) have 4 extra address lines for accessing locations in physical memory. When PAE is enabled, the processor can access up to 64GB of physical RAM. Even though the physical address space is now 36 bits, the size of the logical address space, and the maximum segment size remain unchanged from the i386 at 64TB and 4GB respectively.
- **Using PAE**
- **Logical Memory**
- Despite the fact that a computer may now physically access up to 64GB (236 bytes)of memory, each process still receives at maximum 3GB of *virtual* address space for use. This limit is imposed by the use of 32-bit pointers, and the fact that the kernel reserved the upper 1GB of the 32-bit address space for itself [USE]. Nevertheless, it is possible to work around this limit. For instance, data may be distributed over several processes. In effect, each process could manage a 3GB "chunk" of memory, and data could be passed between processes using standard IPC or shared memory.
- **Physical Memory**
- Physical memory above 1GB is known as the high memory zone [USE]. In earlier 2.2.x releases, this was the maximum amount of physical memory the kernel could use unless it were customized. By 2.2.17 this memory could be accessed by setting the CONFIG_HIGHMEM_4G option which moved the kernel down from 3GB to 2GB. This also cuts down the memory available to any single process however [YAN]. If the CONFIG_HIGHMEM_64G option is selected, then PAE is turned on and the kernel can access up to 64GB of RAM. This support must be compiled into the kernel [CONF]. However, even though your machine may support greater than 4GB of memory, no user process can use more than 4GB of memory at a time (again, the 32-bit pointer limit). Even though pointers are still only 32-bits wide, the kernel can access all of physical memory by changing the values in the top-level directory (pgd) or by changing the base address of the pgd by moving a new value into register CR3[PPR]. (Note that changing CR3 invalidates the entire TLB, and a task switch changes the CR3 [PPR].)

Linux File Systems

- Ext2 – second extended file system
- Ext3 – The ext3 filesystem is a journaling extension to the standard ext2 filesystem on Linux
- Xfs - SGI's XFS filesystem for Linux is a high-performance journaling filesystem that provides rapid recovery from system crashes and the ability to support extremely large disk farms.
- XFS is a full 64-bit filesystem, and thus, as a filesystem, is capable of handling files as large as a million terabytes.
  - 263 = 9 x 1018 = 9 exabytes
  - In future, as the filesystem size limitations of Linux are eliminated XFS will scale to the largest filesystems
  - **Maximum File Size**
  - For Linux 2.4, the maximum accessible file offset is 16TB on 4K page size and 64TB on 16K page size. As Linux moves to 64 bit on block devices layer, file size limit will increase to 9 million terabytes (or the system drive limits).
  - **Maximum Filesystem Size**
  - For Linux 2.4, 2 TB. As Linux moves to 64 bit on block devices layer, filesystem limits will increase.
- LVM – Logical Volume Manager
- VFS – Virtual File System
  - In order to ease the addition of new filesystems into the Linux kernel, a Virtual File System (VFS) layer was developed. The VFS layer was initially written by Chris Provenzano, and later rewritten by Linus Torvalds before it was integrated into the Linux kernel. It is described in The Virtual File System. . The VFS is an indirection layer which handles the file oriented system calls and calls the necessary functions in the physical filesystem code to do the I/O.
- reiser - ReiserFS 3.6.x (the version included as part of Linux 2.4) is designed and developed by Hans Reiser and his team of developers at Namesys. Journaled file system
- NFS

Memory on Linux

- Q: How much memory can a process address in Linux/IA-32? A: IA-32
- Physical memory limits (i.e, total memory recognized by Linux 2.4):
  – Up to 64 GB if hardware supports Intel Physical Address Extensions (PAE) (although the practical limit is ~16GB physical due to certain data-structure design constraints specific to Linux) Up to 4 GB otherwise.
- IA-32 memory limit per process:
  – Up to ~3GB (typically ~2.9 GB) per process (for example, allocated via malloc()) (This limit holds -even- if the hardware supports PAE. Although the aggregate physical memory seen by Linux can be upto 64 GB with PAE, an individual process can only see up to ~2.9 GB).

Application Backup Slides

CAE ISV Itanium 2 Status on HP-UX

15-Jul-03

HP WORLD 2003 Solutions and Technology Conference & Expo

ISV	Application	OS	Port	1st Release	Latest Rel.	Solutions	HP Contact
ABAQUS	ABAQUS Std/Exp	HP-UX			6.3-4	SMP	lee.fisher@hp.com
ACUSIM	AcuSolve	HP-UX			1.4g	cluster/SMP	brian.lowe@hp.com
ADINA R&D	ADINA	HP-UX			8.0	SMP	lee.fisher@hp.com
ADINA R&D	ADINA-AUI	HP-UX			8.0.2	workstation	lee.fisher@hp.com
Altair Engr.	HyperMesh	HP-UX			6.0 rc1	workstation	jeff.wall@hp.com
Altair Engr.	HyperWorks	HP-UX			6.1 Sept	workstation	jeff.wall@hp.com
Altair Engr.	OptiStruct	HP-UX		5	6.0	ws & SMP	jeff.wall@hp.com
ANSYS	ANSYS	HP-UX		6.1	7.0	SMP & w/s	lynnlewis@hp.com
ANSYS	ICEM-CFD/AI*Env	HP-UX		4.3beta	1.0 AI*Env	workstation	lynnlewis@hp.com
ANSYS (CFX)	CFX-5 (solver)	HP-UX		5.5.1	5.6	cluster/SMP	lynnlewis@hp.com
ANSYS (CFX)	TASCflow	HP-UX		2.11.2	2.12	SMP & w/s	lynnlewis@hp.com
AVL	Excite	HP-UX					lynnlewis@hp.com
AVL	FIRE v7	HP-UX			7.3	ws & SMP	lynnlewis@hp.com
AVL	FIRE v8 (SWIFT)	HP-UX	8.1		Q4'03	ws/SMP/clus	lynnlewis@hp.com
BetaCAE Systems	ANSA	HP-UX		pre-rel v11.2.4		workstation	lee.fisher@hp.com
CD adapco Group	PROSTAR	HP-UX			5.2.2	workstation	lynnlewis@hp.com
CD adapco Group	STAR-CD & HPC	HP-UX			3.15A	cluster/SMP	lynnlewis@hp.com
CDH	AMLS	HP-UX			2.0	SMP	jeff.wall@hp.com
CDH	CDH-OPT	HP-UX			2.0	SMP	jeff.wall@hp.com
CEI	EnSight, EnLiten	HP-UX			7.4 & 7.6	workstation	lynnlewis@hp.com
CFD Research	CFD-ACE+	HP-UX		v2003		ws/SMP/clus	jeff.wall@hp.com
Engineous	iSight	HP-UX		7.1	7.2	workstation	lynnlewis@hp.com
ESI Group	OPTRIS	HP-UX					philippe.chaventre@
ESI Group	PAM-CEM	HP-UX			v2002.1	SMP	philippe.chaventre@
ESI Group	PAM-CRASH/SAFE	HP-UX		v2001 SMP	v2002	cluster/SMP	philippe.chaventre@
ESI Group	PAM-FLOW	HP-UX		v2001	v2002.1	cluster/SMP	philippe.chaventre@
ESI Group	PAM-SHOCK/MEDY!	HP-UX		v2001SMP	v2002	cluster/SMP	philippe.chaventre@
ESI Group	PAM-STAMP/FORM	HP-UX		v2001.1 SMP	2G v2003 so	SMP/cluster	philippe.chaventre@
ESI Group	PAM-VIEW	HP-UX			v2002.1	workstation	philippe.chaventre@

Key:
- Green = port underway
- Blue = 1st release version
- Purple = current optimized ver.

Disclaimer: while every effort has been made to verify the details on this table, be sure to check with the local ISV office to confirm. Release status changes every week, so this table will be updated frequently. E-mail the HP contact for assistance on deals.

Note: All listed applications are also on PA-RISC/HP-UX today.

CAE ISV Itanium 2 Status on HP-UX

15-Jul-03

ISV	Application	OS	Port	1st Release	Latest Rel.	Solutions	HP Contact
ESI Group	RAYON	HP-UX					philippe.chaventre@
ESI Group	SYSTUS	HP-UX					philippe.chaventre@
ETA	DYNAFORM	HP-UX			4.0	workstation	lee.fisher@hp.com
ETA	VPG & FEMB	HP-UX			2.5 & 28.0	workstation	lee.fisher@hp.com
Exa	PowerFLOW	HP-UX			3.4 p2a	cluster/SMP	lee.fisher@hp.com
Fluent	FLUENT	HP-UX		6.1.18	6.1.22	cluster/SMP	lee.fisher@hp.com
Fluent	POLYFLOW	HP-UX	3.10		Jul'03	ws & SMP	lee.fisher@hp.com
Fluent	TGrid	HP-UX			3.5	ws & SMP	lee.fisher@hp.com
Gamma Technolog	GT Power	HP-UX					lynnlewis@hp.com
Intes	Permas	HP-UX			v9.0.189	SMP/cluster	lynnlewis@hp.com
LMS	FALANCS	HP-UX	2.13		Dec'03	SMP	lynnlewis@hp.com
LMS	SYSNOISE	HP-UX			5.6	SMP/cluster	lynnlewis@hp.com
LSTC	LS-DYNA	HP-UX		v960.1647	v970SPDP	cluster/SMP	lee.fisher@hp.com
LSTC	LS-PrePost	HP-UX			1.0	workstation	lee.fisher@hp.com
MAGMA	MAGMASOFT	HP-UX	4.2			SMP & w/s	lynnlewis@hp.com
Mecalog	RADIOSS	HP-UX			4.1q SPMD	cluster/SMP	yves.bostroem@hp.
Metacomp	CFD++	HP-UX		3.3.2	3.5.1 Jul'03	cluster/SMP	jeff.wall@hp.com
Moldflow	MPI (serial solver)	HP-UX			4.1	server	lee.fisher@hp.com
MSC.Software	MSC.Akusmod	HP-UX				SMP/cluster	jeff.wall@hp.com
MSC.Software	MSC.Dytran	HP-UX	32-bit ver		dy2004	ws & SMP	jeff.wall@hp.com
MSC.Software	MSC.Marc/Mentat	HP-UX		v2001	v2003	SMP/cluster	jeff.wall@hp.com
MSC.Software	MSC.NASTRAN	HP-UX	70.7 &	v2001.0.4	v2001.0.9	SMP/cluster	jeff.wall@hp.com
MSC.Software	MSC.PATRAN	HP-UX			v2003	workstation	jeff.wall@hp.com
Newmerical	DROP3D	HP-UX			2002R1	SMP	norman.lindsey@hp
Newmerical	Optimesh	HP-UX			4.3	workstation	norman.lindsey@hp
ONERA	elsA	HP-UX					jeff.wall@hp.com
Ricardo	VECTIS	HP-UX		3.6	3.7	cluster/SMP	lynnlewis@hp.com
SAMTECK	SAMCEF	HP-UX			v9.1-05		jeff.wall@hp.com
Software Cradle	SCRYU	HP-UX				SMP	toru.koiso@hp.com
Sofy Technologies	SOFY	HP-UX			3.x beta Jul'03	workstation	jeff.wall@hp.com
TNO	MADYMO	HP-UX		6.0.1p1	6.1	SMP	lynnlewis@hp.com
VNI	IMSL	HP-UX			5		

Key:

Green = port underway

Blue = 1st release version

Purple = current optimized ver.

CAE ISV Itanium 2 Status on LINUX64

15-Jul-03

HP WORLD 2003 Solutions and Technology Conference & Expo

ISV	Application	OS	Port	1st Release	Latest Rel.	Solutions	HP Contact
ABAQUS	ABAQUS Std/Exp	Linux64			6.3-4	SMP	lee.fisher@hp.com
ADINA R&D	ADINA	Linux64			8.0	SMP	lee.fisher@hp.com
ADINA R&D	ADINA-AUI	Linux64			8.0.2	workstation	lee.fisher@hp.com
ANSYS	ANSYS	Linux64		6.3	7.0	SMP & w/s	lynnlewis@hp.com
ANSYS	ICEM-CFD/AI*Env	Linux64		4.3beta	1.0 AI*Env	workstation	lynnlewis@hp.com
ANSYS (CFX)	CFX-5 (solver)	Linux64			beta 5.5.1	cluster/SMP	lynnlewis@hp.com
CD adapco Group	STAR-CD & HPC	Linux64				cluster/SMP	lynnlewis@hp.com
CEI	EnSight, EnLiten	Linux64			7.4 & 7.6	workstation	lynnlewis@hp.com
CFD Research	CFD-ACE+	Linux64	v2003			ws/SMP/clus	jeff.wall@hp.com
ESI Group	PAM-CRASH	Linux64		on demand:	v2002	cluster/SMP	philippe.chaventre@
ESI Group	PAM-FLOW	Linux64		on demand:	v2002.1	cluster/SMP	philippe.chaventre@
Fluent	FIDAP	Linux64			8.7.2	ws & SMP	lee.fisher@hp.com
Fluent	FLUENT	Linux64			6.1.22	cluster/SMP	lee.fisher@hp.com
LSTC	LS-DYNA	Linux64		v960.1106	v970 SP	cluster/SMP	lee.fisher@hp.com
Mecalog	RADIOSS	Linux64	4.1p		Apr'03	cluster/SMP	yves.bostroem@hp.
Metacomp	CFD++	Linux64	3.5.1		Jul'03	cluster/SMP	jeff.wall@hp.com
MSC.Software	MSC.Marc	Linux64				SMP/cluster	jeff.wall@hp.com
MSC.Software	MSC.NASTRAN	Linux64		v2001		SMP/cluster	jeff.wall@hp.com
Ricardo	VECTIS	Linux64	3.6			cluster/SMP	lynnlewis@hp.com

Key:
- Green = port underway
- Blue = 1st release version
- Purple = current optimized ver.

Disclaimer: while every effort has been made to verify the details on this table, be sure to check with the local ISV office to confirm. Release status changes every week, so this table will be updated frequently. E-mail the HP contact for assistance on deals.

Note: All listed applications are also on PA-RISC/HP-UX today.

Application Characteristics

- MSC.NASTRAN
  - NVH
    - dot products and daxbys
    - Two loads per floating point operation
    - Memory bandwidth limited
    - I/O limited
  - Stress
    - Floating point limited
  - Read backwards
- ABAQUS
- AMLS
  - Matrix multiplies ?????dot product
  - Small blocks
  - At the end of the run there can be large I/O

MSC.NASTRAN

DMP is recognized in **SOL 101, 103, 108, 111,** and **112**

Finite Element Analysis

- The use of finite element mathematical methods to simulate the behavior of physical structures, fluids, gasses and magnetic fields when subjected to a wide variety of environmental conditions.
- Uses include:
  - Aerospace- frames, wings, structural design
  - Automotive - body, engines, structural design
  - Civil and architectural engineering - structural integrity for buildings, bridges, tunnels - earthquake proofness
  - Prosthesis design
  - Electronic package design

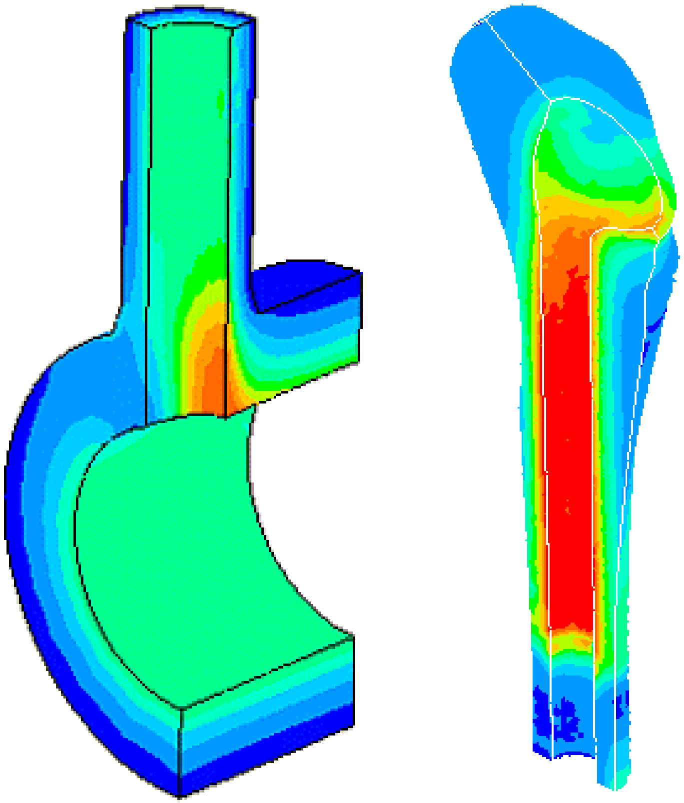

Source: Hibbitt, Karlsson and Sorensen

Cluster Backup slides

System Area Networks:

(Alpha/IA32/IA64)	QSW Elan-3	Myrinet 2000 (rev "C")	100baseT Ethernet	Gigabit Ethernet	SCI Dolphin
Measured Bandwidth (MB/s)	280/305/220	227	8.58	~85 (MPI-CH) ~122 (scaMPI)	142
MPI* Latency (µs)	<5/4.91	<10	115 (100-1000's)	103 (80-1000's)	3.8
Deterministic Transfers	Yes	Yes	No	No	Yes
Barrier/Collective Support	HW/SW	SW	SW	SW	HW/SW
OS Support	Tru64/Linux	Tru64/Linux	Tru64/Linux/HP-UX	Tru64/Linux/HP-UX	Linux
System Management	SC Software w/RMS	Scyld/Scali/ OSCAR/…	Scyld/Scali/ OSCAR/…	Scyld/Scali/ OSCAR/…	Scali Wulfkit
Maximum # nodes	Thousands	Thousands	Thousands	Thousands	Low hundreds
Network Topology	Fat-Tree	Fat-Tree	Star/Layered Tree	Star/Layered Tree	Torus
MPI support	Quadrics/HP	Myricom/MPI-CH	MPI-CH	MPI-CH/ Scali scaMPI	Scali
Average List Price/Node (2-128 Nodes)	$4280	$1941	~$200	$1953	$1682

This data is updated regularly -- see Christian Heiter for updated information. Data as of 2/26/03.

MPI numbers are from the Pallas ping pong User level to User level benchmarks: latency from 0 byte ping-pong, bandwidth for varying sizes.
*Results for QSW from: an AlphaServer ES45 (Quad-processor), 64/66 Elan-3 in an AlphaServer SC configuration.
an Intel P4 Xeon @ 2GHz, dual processor, E7500 chipset, 64/66 Elan-3
*Results for others based on various AlphaServer and Intel IA64 and IA32 cluster configurations
*GbE implementation was with Extreme Networks, non-blocking configuration

See Appendix for information on cost/node as well as comparison for latency and bandwidth with respect to price.

Logical View of a Beowolf Cluster

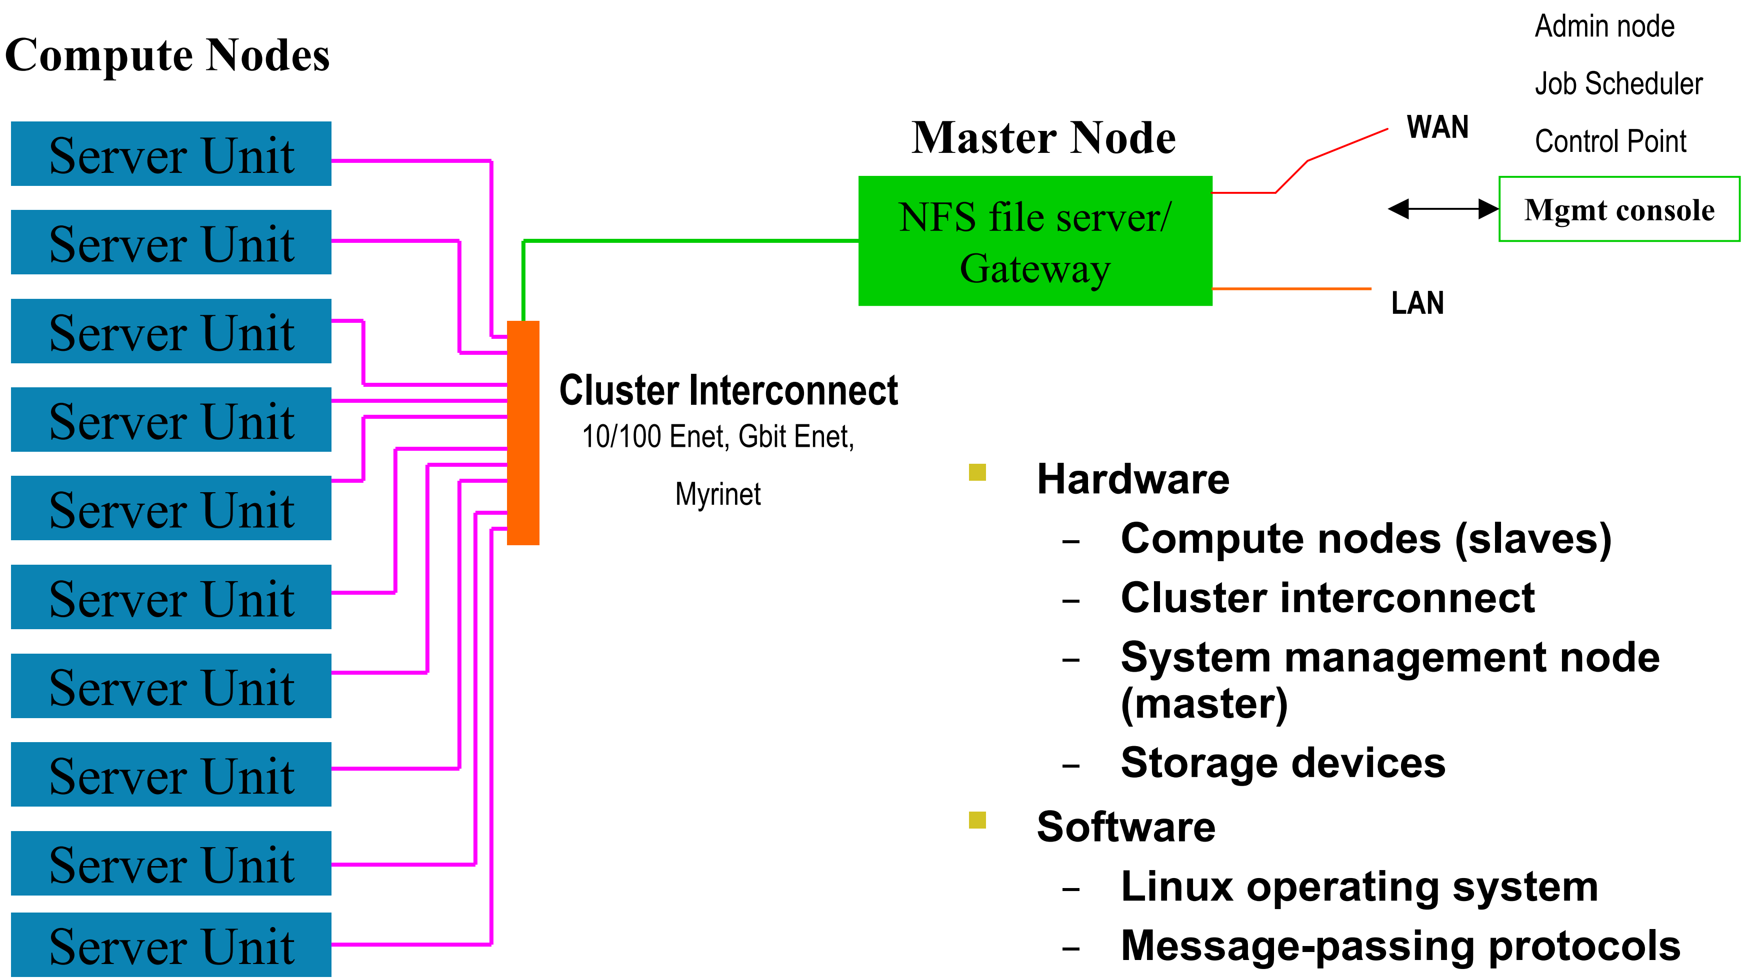

- **Hardware**
  - **Compute nodes (slaves)**
  - **Cluster interconnect**
  - **System management node (master)**
  - **Storage devices**
- **Software**
  - **Linux operating system**
  - **Message-passing protocols**
  - **Parallel application**

HPTC System Architecture Overview