



Platinum Sponsors



Gold Sponsors



Silver Sponsors



Legendary Reliability™



HPC Cluster Solutions



Powerful Cluster Technology



Conference Partner



Media Sponsors



Visit our website for more information about our conference program, exposition and special events.

www.clusterworldexpo.com

ClusterWorld™
Conference & Expo

Explore cutting-edge technology and learn from top experts at ClusterWorld Conference and Expo, the first major event to focus exclusively on cluster technology.

CONFERENCE **JUNE 23-26 2003**
EXPOSITION **JUNE 24-26 2003**
SAN JOSE CONVENTION CENTER
SAN JOSE, CALIFORNIA

www.clusterworldexpo.com



Keynote Speakers

John Picklo, Manager of HPC
DaimlerChrysler



John Reynders, VP for Informatics
Celera Therapeutics



Jack Buur, Principal Research Physicist
Shell Int'l Exploration & Production

REGISTER ONLINE BEFORE MAY 23RD AND SAVE!

What is ClusterWorld?

Introducing ClusterWorld Conference and Expo, the first major event to focus entirely on clustered systems. From research labs to industrial datacenters, clusters are revolutionizing high-performance and high-availability computing.

ClusterWorld's conference program has been created in conjunction with the Linux Clusters Institute and offers something for everyone who works with clusters—whether you're doing pure research, parallel software development, clustered systems administration, or even IT management.

The conference also offers in-depth vertical tracks in fields ranging from bioinformatics to petroleum exploration, CAD, aerospace and automotive engineering, digital content creation, and financial services. Distributed cluster computing is represented as well, with a special focus on emerging Grid computing technologies.

Whatever industry you work in, there is a specific ClusterWorld conference track that will keep you on top of the latest tools and techniques in your field. And the ClusterWorld Expo show floor is the best place to test drive the latest cluster technologies from all the major systems vendors.

If you work with clusters in any capacity, or if you are thinking about deploying cluster technology within your organization, ClusterWorld Conference and Expo is the one event you cannot afford to miss this year.

Register online by May 23rd and take advantage of our money saving Early Bird Conference Package at:

www.clusterworldexpo.com

ClusterWorld Conference and Expo. The power of many.

Event Features

Beowulf Reunion Tour

Thursday, June 26, 2003 2:30pm - 3:30pm

This year marks the 10th anniversary of the Beowulf project, the starting point for Linux PC clustering. The Beowulf Reunion Tour brings together key members of the Beowulf team whose insight and forethought revolutionized clustering.

Don Becker, Tom Sterling and other members of the core team will talk about Beowulf's early beginnings, its successes over the past 10 years and what the next decade may hold.

**Beowulf
Reunion
Tour**

Excellence in Cluster Technology Awards

Wednesday, June 25, 2003 6:15pm - 7:30pm

Clusters are fast becoming the preferred solutions to technical problems requiring low cost, high-performance and high-availability solutions. As this trend continues, there are goods, services and solutions that stand out above the rest. The Excellence in Cluster Technology Awards are given in recognition of outstanding products, services and platforms that advance the role and adoption of clustered systems.

Excellence in
CLUSTER
Technology

GridWars: Parallel Programming Championship

Presented by Engineered Intelligence

Sponsored by HP

Tuesday, June 24th, 12:00pm - 1:00pm

Created to increase interest in parallel programming, GridWars is a highly strategic but simple game where battle programs fight for control over parallel processor cycles. Flex your programming muscle against some of the best cluster developers in the world on the gladiatorial grid.



grid wars II

Visit GridWars.com for information on registration, toolkit downloads, and the rules of the event.

ClusterWorld Challenge

Tuesday, June 24th, 2:30pm - 3:30pm

Cluster experts face off against one another in a quiz show setting at the ClusterWorld Challenge. Hosted by Pete Beckman, the CW Challenge is a combination of Family Feud and post-doctorate-level supercomputing with questions touching on several categories of clustering. Teams compete for prizes and the bragging rights of being ClusterWorld Challenge Champions.

**ClusterWorld
CHALLENGE**

Is your organization interested in participating the ClusterWorld Challenge and going up against some of the world's finest clustering minds?

Cluster Crash Party

Tuesday, June 24th, 6:00pm - 7:00pm

After a hard day of attending conference sessions and checking out the exhibits, it's time for the systems to come down! Join us on the expo floor for our Cluster "Crash" Party. The Cluster Crash is an opportunity for attendees and exhibitors to meet in a relaxed atmosphere over food and beverages. Don't miss this great opportunity to network with your peers or chat with vendor reps!

**Cluster
Crash
Party**

Keynotes



John Picklo

Manager, High Performance Computing, DaimlerChrysler

Mr. Picklo is responsible for systems software and hardware for all of the engineering mainframes and supercomputers at the Chrysler Group. His background includes 25 years of experience working with information technology in various technical and consulting roles. Mr. Picklo's automotive background includes experience designing and managing systems to support computer-aided design at DaimlerChrysler, General Motors, Nissan, and Toyota.



Dr. John Reynders

Vice President, Informatics, Celera Therapeutics

Dr. Reynders is responsible for computational sciences, algorithmics, software engineering, computer science, and knowledge management efforts in support of drug discovery and development at Celera Therapeutics. Previously, Dr. Reynders served as Vice President for Information Systems at Celera Genomics where he was responsible for all supercomputing capabilities, discovery software engineering, and enterprise system infrastructure. Prior to Celera, Dr. Reynders worked at Sun Microsystems, Inc. and Los Alamos National Laboratory where he managed the largest dedicated unclassified super computer in the United States.



Jacobus N. Buur

Principal Research Physicist, Shell International Exploration and Production B.V.

Mr. Buur is responsible for research initiatives in subsurface imaging and exploration for Royal Dutch/Shell via their GameChanger process. He has been instrumental in developing key visualization technologies and has made significant contributions to the adoption of clusters in the petroleum industry with the early installation of large-scale Linux clusters. Mr. Buur has been with the Royal Dutch/Shell family of companies for 20 years.



Dr. Tilak Agerwala

Vice President, Systems, IBM Research

Dr. Agerwala is responsible for developing the next-generation hardware and software technologies for embedded systems, servers, and supercomputers for IBM. Dr. Agerwala's distinguished career includes jointly developing the architectural foundations of the RS6000 and was responsible for the systems architecture and technology strategy of the RS/6000 SP (1992-1997), the most successful parallel computer of all time. Dr. Agerwala received the W. Wallace McDowell Award from the IEEE in 1998 for outstanding contributions to the development of high performance computers.

Advisory Board



Nanette Boden

Executive Vice President, Myricom

Nanette Boden was a member of the group that founded Myricom in 1994. She has served in a number of management positions at Myricom. Her current responsibilities range from oversight of operations to managing key customer and sales-channel relationships.



Jay Clark

Marketing and Business Development, High Performance Computing Group, MSC.Software

Jay Clark was instrumental in the early adoption of clusters into the automotive and aerospace industries. He has been a member of the MSC.Software team for over 13 years and has held positions in Application Engineering and Sales.



Dave Driggers

Founder and Chief Executive Officer, RackSaver

David Driggers is cofounder and CEO of RackSaver, a leading provider of high-density servers and supercomputing clusters. Mr. Driggers is also a member of RackSaver's Board of Directors, as well as being chairman of the company's executive and technology development committees.



Frank L. Gilfeather

Interim Associate Vice Provost for Research, University of New Mexico

Frank L. Gilfeather, is currently Interim Associate Vice Provost for Research at the University of New Mexico with responsibilities involving campus wide research initiatives and planning. He was Executive Director of the HPC Education and Research Center (HPCERC) at UNM from 1993 to 2003.

Rick Herrmann

Industry Manager for High-performance Computing, Intel



Ron Neyland

Director of Cluster Systems Engineering, RLX Technologies

Ron Neyland joined RLX with a primary goal of building a team to deliver Microsoft Windows support for RLX systems. He is currently responsible for RLX's compute cluster solutions.



Pratap Pattnaik

Senior Manager of the Scalable Systems Group, IBM Research Division

Pratap Pattnaik's current research work includes the development and design of computer systems, operating systems and autonomic components. He is also leading the GRID research activities at IBM Research.



Daniel A. Reed

Director, National Center for Supercomputing Applications (NCSA)

Daniel Reed provides strategic direction and vision to NCSA, the National Computational Science Alliance, and the TeraGrid project. He is a respected leader in the national computer science community and among the federal agencies that support research and development.



Dave Rich

Director of HPC Initiatives, AMD

David Rich is a director concentrating on AMD's high performance computing initiatives. Previous to AMD, Mr. Rich was general manager of API NetWorks, a developer of high-speed, I/O interconnect components based on AMD's HyperTransport technology.



Reza Rooholamini

Director of Operating Systems and Cluster Engineering Group, Dell

Reza Rooholamini is responsible for developing the Linux OS and cluster products in the Enterprise Systems Group at Dell. He has over 30 publications in areas of his research interest: HPC, storage systems, HA and interconnects.

PROGRAM COMMITTEE

Donald Becker

Scyld Computing Corporation, USA

Peter Beckman

Argonne National Laboratories, USA

Gianfranco Bilardi

University of Padova, Italy

Jay Boisseau

TACC / UT - Austin, USA

Ron Brightwell

Sandia National Laboratories, USA

Giri Chukkappalli

San Diego Supercomputing Center, USA

Stefano Cozzini

INFM - Sissa, Italy

Cesar De Rose

CPAD - PUCRS / HP, Brazil

Jack Dongarra

University of Tennessee, USA

Thomas Fahringer

University of Vienna, Austria

Wolfgang Gentsch

Sun Microsystems, Inc., USA

Patrick Geoffray

Myricom, USA

David Henry

Linux NetworX, USA

Haeng Jin Jang

KISTI, South Korea

Patricia Kovatch

San Diego Supercomputing Center, USA

Werner Krotz

Vogel - Pallas, Germany

Greg Lindah

Key Technologies, USA

Arthur B. Maccabe

University of New Mexico, USA

Tim Mattson

Intel, USA

Marshall Kirk McKusick

Marshall Kirk McKusick Consultancy, USA

Bart Miller

University of Wisconsin, USA

Ron Minnich

Los Alamos National Laboratories, USA

Bernd Mohr

Research Center Juelich, Germany

Shirley Moore

University of Tennessee, USA

David Morton

MHPCC, USA

J.P. Navarro

Argonne National Laboratories, USA

Henry Neeman

OSCAR - Oklahoma University, USA

Takashi Ohta

IBM Tokyo Research Lab, Japan

Rod Oldehoef

Los Alamos National Laboratories, USA

Jairo Panetta

INPE / CPTEC, Brazil

Michael Pflugmacher

NCSA / UIUC, USA

Christoph Pospiech

IBM, Germany

Jean-Pierre Prost

BM, France

Faisal Saied

NCSA / UIUC, USA

Stephen Scott

ORNL, USA

Nils Smeds

KTH, Sweden

Brian Smith

University of New Mexico, USA

Martin Streicher

Linux Magazine, USA

John Taylor

Quadrics, UK

Jeff Vetter

LLNL, USA

Steering Committee

Robert Ballance

Research Associate Professor, University of New Mexico

Luiz Derose

Research Staff Member, IBM T.J. Watson Research Center

David Klepacki

Senior Staff Member, IBM T.J. Watson Research Center

Becky McGreal

Director Operations, Conference Program LCI & CWCE

Candace Shirley

Program Manager, HPCERC, University of New Mexico

John Towns

Division Director for Scientific Computing, NCSA at University of Illinois

Sponsors as of April 1, 2003

Platinum Sponsors



Gold Sponsors



Silver Sponsors



Conference Partner



Media Sponsors



Exhibitors: SuSE, Cyclades, Dolphin Interconnect

Tutorials

Developed by the Linux Clusters Institute, the ClusterWorld tutorials are hands-on workshops where attendees learn from the real-world experience of cluster experts. The tutorials follow two tracks:

Track 1 (TA): Introduction to Linux Clustering

This track contains tutorials that illustrate starting from the ground up how to design, configure and install a Linux cluster. Topics include:

- Cluster components
- Networking choices
- Hardware alternatives
- System configuration alternatives
- System acceptance testing
- Cluster management tools

Track 2 (TS): The Linux Cluster Software Stack

This track contains tutorials that presents the basic components in the cluster software stack. Topics include:

- Kernel issues
- Compilers
- Libraries
- Tools
- Security
- Scheduling
- Message passing
- Grid components

Monday, June 23, 2003 8:30am -12:00pm

TA1: Programming in the Linux Cluster Environment I: Tools and Program Optimization Techniques

This tutorial presents topics focused on optimizations for cluster applications. Particular attention is given to processor architecture considerations and programming techniques that take advantage of processor architecture. In addition, performance monitoring and analysis techniques and tools will be examined.

Monday, June 23, 2003 8:30am -12:00pm

TS1: Introduction to Linux Clusters: Design Alternatives, Configuration and Installation

In this tutorial, attendees will learn how to design, configure and install a Linux cluster. Topics include choosing cluster components, networking choices, hardware alternatives, system configuration alternatives, system acceptance testing and cluster management tools.

Monday, June 23, 2003 1:30pm - 5:00pm

TA2: Programming in the Linux Cluster Environment II: Parallel Programming with MPI

In this tutorial, attendees will be presented with basic information on programming with the MPI library. Tools for debugging parallel applications will be covered, along with techniques for improving application scalability on Linux clusters.

Monday, June 23, 2003 1:30pm - 5:00pm

TS2: The Linux Cluster Software Stack: Operations from the System Administration Perspective

This tutorial presents the basic components in the cluster software stack: kernel issues, compilers, libraries, tools, security, scheduling, message passing and Grid components. Particular attention is given to the alternatives and tradeoffs in different styles of cluster operations.

Day 1 | Tuesday, June 24, 2003

 Check www.clusterworldexpo.com for updates to the conference program

Time	System	Applications	Cluster Solutions	Bioinformatics	Digital Content Creation/ Visualization/Simulations	Petroleum/ Geophysical Exploration	Automotive & Aerospace Engineering
8:30am - 9:00am	Breakfast						
9:00am - 10:00am	Opening Keynote: John Picklo, DaimlerChrysler						
10:00am - 10:30am	Conference Attendee Break						
10:30am - 11:30am	S1: Simple Linux Utility for Resource Management 10:30am - 11:15am S2: Simple Installation and Administration Tool for Large Clusters 11:15am - 12:00pm	A1: Large Scale Parallel Reservoir Simulations on a Linux PC Cluster 10:30am-11:15am A2: Scalable Performance of FLUENT on NCSA IA-32 Linux Cluster 11:15am-12:00pm	C1: Building the TeraGrid: The World's Largest Grid, Fastest Linux Cluster, and Fastest Optical Network Dedicated to Open Science	B1: Running BLAST on a Linux Cluster	D1: The Current State of Numerical Weather Prediction on Cluster Technology - What is Needed to Break the 25% Efficiency Barrier?	P1: Exploring the Earth's Subsurface with Itanium 2 Linux Clusters	E1: Cluster Computing in Space Applications
11:30am - 1:15pm	Visit Exhibits Lunch GridWars Parallel Programming Championship 12:00pm-1:00pm						
1:15pm - 2:15pm		S3: The Space Simulator	C2: Building x86-64 Applications for AMD Opteron HPC Clusters	B2: Biobrew Linux: A Linux Cluster Distribution For Bioinformatics	D2: Building and Using Tiled Display Walls	P2: Scalability Considerations for Compute Intensive Applications	E2: Full Vehicle Dynamic Analysis Using Automated Component Modal Synthesis
2:30pm - 3:30pm	ClusterWorld Challenge hosted by Pete Beckman						
3:30pm - 4:00pm	Conference Attendee Break						
4:00pm - 5:00pm	S4: A Middleware-Level Parallel Transfer Technique Over Multiple Network Interfaces	A3: Moore's Law and Cluster Computing: When Moore Is Not Enough	C3: Tools for Optimizing HPC Applications on Intel Clusters	B3: Terascale Linux Clusters: Supercomputing Solutions for Life Sciences	D3: Real-time Visualization of Cluster Networks	P3: Parallel Reservoir Simulation on Intel Xeon HPC Clusters	E3: Using Clusters to Deliver Turn Key CFD Solutions
6:00pm - 7:00pm	Cluster Crash Party						
7:15pm - 9:00pm	Birds-of-a-Feather Meetings						

Day 2 | Wednesday, June 25, 2003

Time	System	Applications	Cluster Solutions	Bioinformatics	Digital Content Creation/ Visualization/Simulations	Petroleum/ Geophysical Exploration	Automotive & Aerospace Engineering
8:30am - 9:00am	Breakfast						
9:00am - 10:00am	Keynote: Jacobus N. Buur, Shell International Exploration and Production B.V.						
10:00am - 10:30am	Conference Attendee Break						
10:30am - 11:30am	S5: The Cluster Integration Toolkit (CIT) 10:30am - 11:15am S6: Scalable C3 Power Tools 11:15am - 12:00pm	A4: Cooperative Caching in Linux Clusters 10:30am - 11:15am A5: Object storage: scalable bandwidth for HPC clusters 11:15am - 12:00pm	C4: The Ultra-Scalable HPTC Lustre Filesystem	B4: Blade Servers for Genomic Research	D4: HPC and HA Clustering for Online Gaming	P4: Geoscience visualization and seismic processing clusters: collaboration and integration issues	E4: LS-DYNA: CAE Simulation Software on Linux Clusters
11:30am - 1:15pm	Visit Exhibits Lunch						
1:15pm - 2:15pm		S7: Full Circle: Simulating Linux Clusters on Linux Clusters	C5: Building the World's Most Powerful Cluster		D5: Large Scale Scientific Visualization on PC Clusters	P5: Case Study: Deploying Large-scale Seismic Processing Clusters at CGG	E5: Linux Clusters in the German Automotive Industry
2:30pm - 3:30pm	Keynote: Tilak Agerwala, IBM Research						
3:30pm - 4:00pm	Conference Attendee Break						
4:00pm - 5:00pm	S9: Scheduling for Improved Write Performance in a Parallel Virtual File System (CEFT-PVFS) 4:45pm - 5:30pm	A6: Analyzing Cluster Log Files Using Logsurfer 4:00pm - 4:45pm A7: Performance Evaluation of Load Sharing Policies with PANTS on Beowulf Clusters 4:45pm - 5:30pm	C6: Driving Cluster/Grid Technologies in HPC	B6: High Performance Mathematical Libraries for Itanium 2 Clusters	D6: The Power of Simulations: Predicting Emergent Behavior	P6: Grid Computing In The Energy Industry	E6: Improving Multi-site/Multi-departmental Cluster Systems Through Data Grids in the Automotive and Aerospace Industries
5:00pm - 6:00pm	Visit Exhibits						
6:15pm - 7:30pm	Excellence in Cluster Technology Awards						

Day 3 | Thursday, June 26, 2003

Time	System	Applications	Cluster Solutions	Bioinformatics	Digital Content Creation/ Visualization/Simulations	Petroleum/ Geophysical Exploration	Automotive & Aerospace Engineering
8:30am - 9:00am	Breakfast						
9:00am - 10:00am	Keynote: John Reynders, Celera Therapeutics						
10:00am - 10:30am	Conference Attendee Break						
10:30am - 11:30am	S10: Achieving Order through CHAOS: The LLNL HPC Cluster Experience 10:30am-11:15am S11: Supercomputing Center Management Using AIRS 11:15am - 12:00pm	A8: On the Numeric Efficiency of C++ Packages in Scientific Computing 10:30am-11:15am A9: Benchmarking I/O Solutions for Clusters 11:15am-12:00pm	C7: Emerging Trends in Data Center Powering and Cooling	B7: Parallel Computational Biology Tools and Applications for Windows Clusters	D7: The Use of Clusters for Engineering Simulation	P7: Drilling in the Digital Oil Field: High Pay-offs from Linux Clusters	E7: Scrutinizing CFD Performance on Multiple Linux Cluster Architectures
11:30am - 1:15pm	Visit Exhibits Lunch						
1:15pm - 2:15pm		A10: The Design, Implementation, and Evaluation of mpiBLAST	C8: The Virtual Environment and Its Impact on IT Infrastructure	B8: Building Software for High Performance Informatics and Chemistry	D8: NEESgrid: Virtual Collaboratory for Earthquake Engineering and Simulation		E8: Managing CAE Simulation Workload in Cluster Environments
2:30pm - 3:30pm	Beowulf Reunion Tour						

Conference Sessions

SYSTEMS TRACK

Tuesday, June 24, 2003

10:30am -11:15am

S1: Simple Linux Utility for Resource Management (SLURM)

Morris Jette, Lawrence Livermore National Laboratory

SLURM is an open source, fault-tolerant and highly scalable cluster management and job scheduling system for Linux clusters of thousands of nodes. Components include machine status, partition management, scheduling and stream copy modules. This session presents an overview of the SLURM architecture and functionality.

Tuesday, June 24, 2003

11:15am-12:00pm

S2: Simple Installation and Administration Tool for Large Clusters

Tomoyuki Hiroyasu, Doshisha University

The installation and configuration of clusters with many nodes is difficult due to the large amount of time and knowledge required to fully complete the task. To solve this problem a simple installation and administration tool, "Doshisha Cluster Auto Setup Tool: DCAST," has been developed. Targeted at Linux, it supports both diskless and diskfull clusters, requires no interaction during install, boots slave nodes over the network and changes to configuration are propagated to all nodes.

Tuesday, June 24, 2003

1:15pm-2:15pm

S3: The Space Simulator

Michael S. Warren, Los Alamos National Laboratory

The Space Simulator is a 294 processor Beowulf cluster with a peak performance near 1.5 Teraflops. It achieved Linpack performance of 665.1 Gflops on 288 processors, making it the 85th fastest computer in the world. The Space Simulator Cluster is dedicated to performing computational astrophysics simulations in the Theoretical Astrophysics group (T6) at Los Alamos National Laboratory. This case study will outline the design drivers, software and applications applied to the system and lessons learned in building this cluster.

Tuesday, June 24, 2003

4:00pm -5:00pm

S4: Middleware-Level Parallel Transfer Technique Over Multiple Network Interfaces

Nader Mohamed, University of Nebraska, Lincoln

Network middleware is a software layer that provides abstract network APIs to hide the

lowlevel technical details from users. Existing network middleware support single network interface and link message transfers. In this session, we describe a middleware level parallel transfer technique that utilizes multiple network interface units that may be connected through multiple networks. It operates on any reliable transport protocol such as TCP and transparently provides an expandable high bandwidth solution that reduces message transfer time, provides fault tolerance and facilitates dynamic load balancing between the underlying multiple networks. The experimental evaluation displayed a peak performance of 187Mbps on two fast Ethernet networks.

Wednesday, June 25, 2003

10:30am -11:15am

S5: The Cluster Integration Toolkit (CIT)

James H. Laros III, Sandia National Labs

The Cluster Integration Toolkit is an extensible, portable, scalable cluster management software architecture for a variety of systems. It has been successfully used to integrate and support a number of clusters at Sandia National Labs and several other sites, the largest of which is 1861 nodes. This session will discuss the goals of the project and how they were achieved. The installation process will be described and common tasks for cluster implementation and support will be demonstrated.

Wednesday, June 25, 2003

11:15am -12:00pm

S6: Scalable C3 Power Tools

Stephen Scott, Oak Ridge National Laboratory

With the growth of the typical cluster reaching 512 and more compute nodes, it is apparent that cluster tools must begin to reach toward the 1000's of nodes in scalability. Version 3.2 of the C3 tools has started stretching the Single System Illusion concept into the realm of 1000's of compute nodes by actually improving performance on larger clusters. This session is a discussion of how this was implemented and how to use this new version of C3 and also presents some results comparing the latest release with prior versions of C3.

Wednesday, June 25, 2003

1:15pm-2:15pm

S7: Full Circle: Simulating Linux Clusters on Linux Clusters

Jose Moreira, IBM Thomas J. Watson Research Center

BGLsim is a complete system simulator for parallel machines allowing users to develop,

test and run the same code that will be used in a real system. It is currently being used in hardware validation and software development for the BlueGene/L cellular architecture machine. BGLsim is capable of functionally simulating multiple nodes of this machine operating in parallel. It simulates instruction execution in each node and the communication that happens between nodes. To illustrate the capabilities of BGLsim, experiments running the NAS Parallel Benchmark IS on a simulated BlueGene/L machine are described.

Wednesday, June 25, 2003

4:45pm -5:30pm

S9: Scheduling for Improved Write Performance in a Parallel Virtual File System (CEFT-PVFS)

Yifeng Zhu, University of Nebraska, Lincoln

This session will demonstrate that all the disks on the nodes of a cluster can be connected together through CEFTPVFS, an RAID10 style parallel file system for Linux system, to provide a GBytes/sec parallel I/O performance, without any additional cost. To improve the overall I/O performance, I/O requests can be scheduled on a less loaded node in each mirroring pair, thus making more informed scheduling decisions. Based on the heuristic rules we found from the experimental results, a scheduling algorithm for dynamic load-balancing has been developed that significantly improves the overall performance.

Thursday, June 26, 2003

10:30am -11:15am

S10: Achieving Order through CHAOS: The LLNL HPC Cluster Experience

Robin Goldstone, Lawrence Livermore National Laboratory

For the past several years, Lawrence Livermore National Laboratory (LLNL) has invested significant effort in the deployment of large High Performance Computing (HPC) Linux clusters. After deploying two modest sized clusters (88 nodes and 128 nodes) in early 2002, efforts progressed to the deployment of the Multi-programmatic Capability Resource (MCR, 1154 nodes) in fall 2002 and ASCI Linux Cluster (ALC, 962 nodes) in early 2003. Through these efforts, LLNL has developed expertise in a number of areas related to the design, deployment and management of large Linux clusters. In this session LLNL will present their experiences, including challenges encountered and lessons learned.

Thursday, June 26, 2003

11:15am-12:00pm

S11: Supercomputing Center Management Using AIRS

Robert A. Ballance, University of New Mexico

Running a large university supercomputing center teaches many lessons, including the need to centralize data collection and analysis, automate system administration functions, and enable users to manage their own projects. Albuquerque Integrated Reporting System (AIRS), a centralized, web-enabled application capable of user and project administration across multiple clusters and reporting against both active and historical data, evolved in response to these pressures.

APPLICATIONS TRACK

Tuesday, June 24, 2003

10:30am -11:15am

A1: Large Scale Parallel Reservoir Simulations on a Linux PC Cluster

Walid A. Habiballah, Petroleum Engineering Application Services Department

Numerical simulation is an important tool used by engineers to develop production strategies and enhance hydrocarbon recovery from reservoirs. Demand for large scale reservoir simulations is increasing as engineers want to study larger and more complex models. In this study, we evaluate a state of the art PC cluster and available software tools for production simulations of large reservoir models. We discuss some of our findings and issues related to large scale parallel reservoir simulations and present performance comparisons between a Pentium IV Linux PC cluster and an IBM SP Nighthawk supercomputer.

Tuesday June 24, 2003

11:15am-12:00pm

A2: Scalable Performance of FLUENT on NCSA IA-32 Linux Cluster

Wai Yip Kwok, National Center for Supercomputing Applications

FLUENT, a leading industrial computational fluid dynamics (CFD) software, has been ported to the NCSA IA-32 Linux cluster. For this study, the scalable performance of FLUENT is benchmarked with two engineering problems from Caterpillar, Inc. and Fluent, Inc with a maximum of 64 processors to accommodate up to 10 million cells. This session will outline the impacts of different interconnects on simulation performance. Using Myrinet

Conference Sessions

interconnects, the Linux cluster computes more than 2.5 times faster than an SGI Origin2000 supercomputer at NCSA. A performance increase of seven times is observed when 32 processors are used instead of two.

Tuesday June 24, 2003

4:00pm -5:00pm

A3: Moore's Law and Cluster Computing: When Moore Is Not Enough

Greg Lindahl, Key Research, Inc.

Linux cluster builders have become accustomed to continuous improvement of cluster building blocks: each year, CPUs get faster, disks get bigger, memory bandwidth rises and networks get cheaper and faster. These improvements are often seen as the inevitable march of progress, driven by the commodity market and Moore's Law. This session will revisit Moore's famous law in detail to determine if it adequately predicts an environment ripe for commodity cluster computing.

Wednesday June 25, 2003

10:30am -11:15am

A4: Cooperative Caching in Linux Clusters

Ying Xu, University of California Riverside

Operating systems used in most Linux clusters only manage memory locally without cooperating with other nodes in the system. This can create states where a node within the cluster may be short of memory while idle memory in other nodes is wasted. This session attempts to solve the problem of how to improve the cluster operating system to support the use of cluster-wide memory as a global distributed resource. Presented will be a description of a cooperative caching scheme for caching files in the cluster-wide memory and corresponding changes in Linux kernel memory management to support it.

Wednesday June 25, 2003

11:15am -12:00pm

A5: Object storage: Scalable Bandwidth for HPC Clusters

Garth Gibson

This session describes the Object Storage Architecture solution for cost-effective, high bandwidth storage in HPC environments. It addresses the unique problems of storage intensive computations in very large clusters, suggesting that a shared file system with out-of-band metadata management is needed to achieve the required bandwidth. The session further argues that for excellent data reliability, storage protection needs to be supported on the data path and it recommends the higher-level semantics of object-based, rather than

block-based, storage for scalable performance, data reliability and efficient sharing.

Wednesday June 25, 2003

4:00pm-4:45pm

A6: Analyzing Cluster Log Files Using Logsurfer

James Prewett, University of New Mexico

Logsurfer is a log analysis tool that simplifies maintaining a cluster by aiding identification and resolution of system issues. This session will outline several examples of using Logsurfer in a cluster environment. Examples range from finding the traces of a complex exploitation of a service to determining which of a set of nodes have problems rebooting. Attendees will learn to configure Logsurfer to meet the particular needs of their environment.

Wednesday June 25, 2003

4:45pm -5:30pm

A7: Performance of Load Sharing Policies with PANTS on Beowulf Clusters

James Nichols, Worcester Polytechnic Institute

Powerful, low-cost clusters of personal computers such as Beowulf clusters, have fueled the potential for widespread distributed computation. While these Beowulf clusters typically have software that facilitates development of distributed applications, there is still a need for effective distributed computation that is transparent to the application programmer.

Thursday, June 26, 2003

10:30am -11:15am

A8: On the Numeric Efficiency of C++ Packages in Scientific Computing

Ulisses Mello, IBM TJ Watson Research Center Object-Oriented Programming (OOP) has proven to be a useful paradigm for programming complex models. In spite of recent interest in expressing OOP paradigms in languages such as FORTRAN90, C++ is the dominant OO language in scientific computing, despite its complexity. Barton & Nackman advocated C++ as a replacement for FORTRAN in engineering and scientific computing due to its availability, portability, efficiency, correctness and generality. These authors used OOP for code reorganization of LAPACK (Linear Algebra PACKage), and they were able to group and wrap over 250 FORTRAN routines into much smaller set of classes, which expressed the common structure of LAPACK.

Thursday, June 26, 2003

11:15am-12:00pm

A9: Benchmarking I/O Solutions for Clusters

Stefano Cozzini, Democritos INFM National Simulation Centre

Clustered Systems offer many advantages for demanding scientific applications: they can deal with massive CPU-bound requirements and allow the distribution of RAM among many nodes. However, many scientific applications process massive amounts of data and therefore require high performance, distributed storage next to parallel I/O. This session will discuss present-day I/O cluster solutions based on Bonnie performance benchmarking for a variety of popular systems.

Thursday, June 26, 2003

1:15pm-2:15pm

A10: The Design, Implementation, and Evaluation of mpiBLAST

Aaron E. Darling, University of Wisconsin-Madison

mpiBLAST is an Open Source parallelization of BLAST that achieves superlinear speed-up by segmenting a BLAST database and then having each node in a computational cluster search a unique portion of the database. Database segmentation permits each node to search a smaller portion of the database, eliminating disk I/O and vastly improving BLAST performance. Because database segmentation does not create heavy communication demands, BLAST users can take advantage of low-cost and efficient Linux cluster architectures such as the bladed Beowulf. In addition to this presentation of the software architecture of mpiBLAST, there will be a detailed performance analysis of mpiBLAST to demonstrate its scalability.

CLUSTER SOLUTIONS TRACK

Tuesday, June 24, 2003

10:30am -11:30am

C1: Building the TeraGrid

Pete Beckman, Argonne National Laboratory

The TeraGrid is one of the most ambitious collaborative grid projects ever undertaken. The building blocks for the \$88 million National Science Foundation funded project include mammoth computational resources, ultra-fast fiberoptic networks linking NCSA, SDSC, CalTech Argonne and PSC and a software "grid hosting environment." Together, they will form an environment that makes developing cluster-based, grid-enabled scientific applications easy. This presentation will provide an overview of the project, the bleeding edge technologies used to bring clusters and grids to the scientific community and an update on current status and results.

Tuesday, June 24, 2003

1:15pm -2:15pm

C2: Building x86-64 Applications for AMD Opteron HPC Clusters

Richard Brunner, AMD

This session will provide an overview of AMD Opteron architecture and a real world look at HPC applications based on the AMD Opteron processor. The session will focus on LINUX SMP and NUMA applications with an overview of the AMD Opteron family of 64-bit processors and associated support ICs from AMD. The applications session will also demonstrate how to apply HyperTransport Technology with the vast ecosystem of support ICs, which are available today and in the months to come.

Tuesday, June 24, 2003

4:00pm -5:00pm

C3: Tools for Optimizing HPC Applications on Intel Clusters

Don Gunning, Intel

The Intel software research lab is involved in several projects related to the development and deployment of HPC software on Intel based clusters. This discussion will focus on the work Intel is doing in parallel/concurrent computing within a single job or task, the development, debugging and tuning multi-threaded applications, in addition to deploying MPI (and mixed MPI/threaded) applications and Extending OpenMP to execute across clusters. This discussion will also touch on ideas for maximum messaging performance on the interconnect while maximizing application performance on the node.

Conference Sessions

Wednesday, June 25, 2003

10:30am -11:30am

C4: The Ultra-Scalable HPTC Lustre Filesystem
Kent Koeninger, HP

The Lustre filesystem is designed to provide a coherent-scalable shared filesystem that can serve thousands of Linux client nodes, delivering extremely high-bandwidth parallel-filesystem access to many terabytes of storage. This talk will describe how the Lustre filesystem will be used in scalable-HPTC-Linux systems to combine the flexibility, scalability and manageability of NAS systems with the performance of SAN systems. The Lustre development effort is an open source project with initial release target in 2003.

Wednesday, June 25, 2003

1:15pm -2:15pm

C5: Building the World's Most Powerful Cluster

Kim Clark, Linux NetworX

In 2002, Linux NetworX built the MCR cluster housed at Lawrence Livermore National Laboratory. It is currently the largest cluster in the world with a theoretical peak of 11.2 Tflops and, with more than 1,000 nodes to manage and monitor, ranks as the fifth largest supercomputer in the world. The unique challenges involved in building and configuring such a massive system and what was learned from this experience will be discussed. Attendees will learn how to apply aspects of the LLNL system to their own smaller system to enhance cluster performance and reliability.

Wednesday, June 25, 2003

4:00pm -5:00pm

C6: Driving Cluster/Grid Technologies in HPC
David Barkai, Intel

High performance computing has undergone a metamorphosis in the last 15-20 years. The impact of changes to the industry and user community will be reviewed. Discussed will be a summary of the building blocks as a set of components built upon enabling technologies, while highlighting the gaps and challenges as cluster computing ramps up and grid computing continues to develop.

Thursday, June 26, 2003

10:30am -11:30am

C7: Emerging Trends in Data Center Powering and Cooling

Wahib Nawabi, APC

Traditional data center architecture approaches force enterprises to build out to full capacity from day one, yet one hundred percent utilization of the designed capacity is seldom reached.

This results in long deployment schedules, millions of dollars of unrecoverable up-front capital investments and the maintenance of expensive service contracts on under-utilized infrastructure. APC's PowerStruXure offers an on-demand solution that accelerates speed of deployment and allows you to invest in a data center solution that is sufficient to meet today's demands, rather than an uncertain estimate of future capacity.

Thursday June 26, 2003

1:15pm -2:15pm

C8: The Virtual Environment and Its Impact on IT Infrastructure

Daniel Kusnetzky, IDC

IDC has been examining the evolution of the virtual environment for quite a number of years. This session will examine IDC's definition of the virtual environment, its roots in techniques developed in the late 1970s, and how Windows, Unix and Linux can be deployed as platforms in the virtual environment. Dan Kusnetzky, IDC's Vice President of System Software, will present the drivers for virtual environment software adoption and project how the virtual environment will impact the overall IT infrastructure in the coming years.

BIOINFORMATICS TRACK

Tuesday, June 24, 2003

10:30am -11:30am

B1: Running BLAST on a Linux Cluster
Ray Hookway, HP

Everyone knows that Blast is an example of an embarrassingly parallel application, i.e., an application that will run well on a cluster. Conceptually, one breaks up a query against a database into several queries against subsets of the database and distributes the resulting jobs across the nodes of the cluster. However, it is not obvious how to go about doing this. The talk will begin with a brief review of how Blast works and then will explore factors that effect the performance of Blast running on a single system. Final focus will be on the answer to the question "How to run Blast on a cluster?"

Tuesday, June 24, 2003

1:15pm -2:15pm

B2: Biobrew Linux: A Linux Cluster Distribution For Bioinformatics

Glen Otero, Callident

Biobrew Linux is the first known attempt at creating and freely distributing an easy-to-use clustering software package designed for bioinformaticists. With support for both IA32 and IA64 platforms, Biobrew is a Linux distribu-

tion that combines the NPACI Rocks cluster software with several popular Open Source bioinformatics software tools like BLAST, HMMER, ClustalW and BioPerl. The result is a Linux distribution that can be used to install a workstation or a Beowulf cluster for bioinformatics analyses.

Tuesday, June 24, 2003

4:00pm -5:00pm

B3: Terascale Linux Clusters: Supercomputing Solutions for the Life Sciences

Bruce Ling, Tularik

At Tularik, a biotechnology company specializing in drug discovery and development using gene regulation, informatics has become essential for the process of genomics-based drug discovery. With the explosion of the genomic data and lead discovery screening data points, a powerful computing environment becomes a must in order to boost B&D productivity. By deploying a 150-processor cluster, Tularik has successfully managed millions of data points, coming from Assay-Development, High-Throughput-Screening (HTS), Structure-Activity-Relationship (SAR), Lead-Optimization and Micro-Array to speed its R&D productivity and decision making processes.

Wednesday, June 25, 2003

10:30am -11:30am

B4: Blade Servers for Genomic Research
Ron Neyland, RLX Technologies

Clusters based on industry standard hardware and software have become the most widely used tools for performing genomic processing and analysis. While providing many benefits such as outstanding price/performance, they also introduce a new set of problems. This session will address how blade servers provide a compute cluster platform that delivers the compute power required for genomic research, while minimizing many of the problems. Real world examples of clusters running many of the widely used genomic applications will be presented, along with tips and tools for managing the cluster environment.

Wednesday, June 25, 2003

4:00pm -5:00pm

B6: High Performance Mathematical Libraries for Itanium 2 Clusters

Hsin-Ying Lin, HP

HP's Mathematical LIBrary (HP MLIB) provides a user-friendly interface using standard definitions of public domain software and enables users to access the power of high

performance computing. HP MLIB fully exploits the architecture of the processor and achieves optimal performance on Itanium 2. HP MLIB has been used by high performance computing customers for over 15 years. This session will provide a brief overview of relevant architectural features and depict how these features have been used to design high-level algorithms. The performance of some of the key components in HP MLIB on Itanium 2 clusters will be discussed: i.e. matrix multiplication, ScaLAPACK and SuperLU_DIST.

Thursday, June 26, 2003

10:30am -11:30am

B7: Parallel Computational Biology Tools and Applications for Windows Clusters

Jaroslav Pillardy, Cornell Theory Center

Using massively parallel programs for data analysis is the most popular way of dealing with the enormous amounts of data produced in molecular biology research. Several computational biology tools for Microsoft Windows clusters of different levels of complexity, available at the Computational Biology Service Unit at the Cornell Theory Center, will be discussed. All of the tools follow a master-worker approach using MPI communications. The simplest tools - tools that are very important to biologists - are standard sequence-based data mining tools such as BLAST and HMMER. More sophisticated is the structure-based (threading) protein annotation algorithm LOOPP.

Thursday, June 26, 2003

1:15pm -2:15pm

B8: Building Software for High Performance Informatics and Chemistry

Joseph Landman, Scalable Informatics LLC

Given the growth rate of life science data sets, analysis applications designed for single machines with shared memory and one or more CPUs quickly leads to a performance bottleneck. Clusters and Grids represent a potential solution to this bottleneck but only when applications are properly designed to make full use of the resources available. In this session we will look at the hard realities of building software for the informatics industry, including: problems with running legacy software on clusters, how to make efficient use of clusters, for both the cluster and the user, and making life science informatics and chemistry applications scale well on clustered systems.

Conference Sessions

DIGITAL CONTENT CREATION / VISUALIZATION / SIMULATIONS

Tuesday, June 24, 2003

10:30am -11:30am

D1: The Current State of Numerical Weather Prediction on Cluster Technology
Dan Weber, Center for the Analysis and Prediction of Storms

This session will look in depth at the current state of weather prediction and the many challenges it faces. The talk will examine the computational needs (teraflops) of a robust numerical weather prediction (NWP) system at thunderstorm scale and review NWP performance on current computer technology. A review of current models will be addressed, as well as the roadblocks associated with clusters. Finally, will be proposal for a complete shift in the way systems of equations are solved on scalar technology in order to break the 25% efficiency ceiling.

Tuesday, June 24, 2003

1:15pm-2:15pm

D2: Building and Using Tiled Display Walls
Paul John Rajlich, National Center for Supercomputing Applications

Tiled display walls provide a large-format environment for presenting very high-resolution visualizations by tiling together the output from a collection of projectors. Projectors are driven by a Linux cluster augmented with high-performance graphics accelerator cards and costs are controlled by using commodity projectors and low-cost PCs. Tiled walls must face a number of challenges, such as, aligning the projectors so that the output of adjacent tiles align to create a seamless image. This session will discuss the Alliance Display Wall-in-a-Box effort; a distribution of related Open Source software packages that reduces the setup and maintenance of complex high-end display systems.

Tuesday, June 24, 2003

4:00pm-5:00pm

D3: Real-time Visualization of Cluster Networks

Tom Caudell, University of New Mexico

The real-time visualization of cluster networks provides a number of benefits to administrators and developers in search of performance bottlenecks. Real-time visuals provide early warning of real problems in network traffic as well as provide clear indication of potential problems before they occur. However, real-time network visualization is a remarkably difficult project. This session will discuss a number of the technical hurdles involved in building a visualization

system that will scale with increased performance. Using network visualization, organizations can design applications that take better advantage of network traffic, avoiding bottlenecks, and administrators can make informed decisions on scheduling that lead a cluster toward optimal performance.

Wednesday, June 25, 2003

10:30am -11:30am

D4: HPC and HA Clustering for Online Gaming

Jesper Jensen, William Pentoja, Matthew Elkourie, SCI

SCI, the company who developed and supports the backend for the Department of Defense's America's Army game, will deliver a case study on deploying gaming clusters for the DoD — and other game titles — and give an overview of where large-scale game technology is and where it is going. With technology capable of pushing an average of 1.35 teraflops per cabinet space and leveraging multiple transit carriers, SCI clusters deliver both the HPC and HA required to support a massive gaming audience. This discussion will touch on solutions for 32 bit and next-generation 64 bit architectures both in place and under development.

Wednesday, June 25, 2003

1:15pm -2:15pm

D5: Large Scale Scientific Visualization on PC Clusters

Brian Wylie, Sandia National Laboratories

This session covers the use of PC clusters with commodity graphics cards as high-performance scientific visualization platforms. A cluster of PC nodes, in which many or all of the nodes have 3D hardware accelerators, is an attractive approach to building a scalable graphics system. The main challenge in using cluster-based graphics systems is the difficulty of realizing the full aggregate performance of all the individual graphics accelerators. Topics covered will include parallel geometric rendering, parallel volume rendering, data distribution approaches and novel techniques for utilizing graphics processors.

Wednesday June 25, 2003

4:00pm -5:00pm

D6: The Power of Simulations: Predicting Emergent Behavior
Cameron Hunt

Commodity hardware and Open Source software like Linux have made distributed computing through clusters and Grids the most cost-effective method of increasing computer power. By combining the ubiquity of distributed comput-

ing resources with a semantic approach to distributed and federated data sources, organizations can perform simulations that create predictive emergent behavior. This session will examine modern predictive efforts that focus on brute force calculation and future trends in simulation with emphasis on the distribution of both data and computational resources.

Thursday, June 26, 2003

10:30am-11:30am

D7: The Use of Clusters for Engineering Simulation

Lynn Lewis, HP

Clusters allow the use of advanced mathematical techniques for optimization, changing the way engineers arrive at cost effective, safe designs. Without inexpensive clusters, engineers at automotive manufacturers could not do 1000's of crash test simulations integrated with the initial design stage nor test for structural integrity much less manufacturability within weeks. This session will examine in detail how, over the previous decade, Unix and lately Linux clusters have found use in commercial cash and fluid dynamics simulations, changing the way cars and aircraft are designed and built.

Thursday June 25, 2003

1:15pm-2:15pm

D8: NEESgrid: Virtual Collaboratory for Earthquake Engineering and Simulation

Tom Prudhomme, National Center for Supercomputing Applications

NEESgrid will link earthquake engineering researchers across the U.S. with leading-edge computing resources and research and testing facilities, allowing teams to plan, perform, and publish their research. Via both Telepresence and other collaboration technologies, research teams are able to work remotely on experimental trials and simulations. This session will examine how NEESgrid, through the shared resources of Grid technology, will bring together information technologists and engineers in a way that will revolutionize earthquake engineering, research and simulation.

PETROLEUM TRACK

Tuesday June 24, 2003

10:30am -11:30am

P1: Exploring the Earth's Subsurface with Itanium 2 Linux Clusters
Keith Gray

This case study of an Itanium 2 processor architecture and Linux cluster technology for seismic imaging and migration imperative project, allowed a major Oil & Gas company to

reduce their cost for this high-end infrastructure by one-half while increasing performance by 3X, and in some cases exceeding this expectation by 5X. The environment includes 1024 processors (4-way HP rx5670 servers x 256 servers) with 8.2 Terabytes (32GB per rx5670 server) of memory and operates at over 4Teraflops peak performance.

Tuesday June 24, 2003

1:15pm -2:15pm

P2: Scalability Considerations for Compute Intensive Applications

Christian Tanasescu, SGI

This session investigates the scalability, architectural requirements and performance characteristics of some of the most widely used compute intensive applications in the scientific and engineering communities. Seismic Processing and Reservoir Simulation (SPR) applications generally consume data read from memory and have to load continuous new data. As a result, to keep the floating point (FP) units busy, these applications require computer architectures with high memory bandwidth, mainly due to the data addressing patterns and heavy I/O activities. We will also introduce BandeLa, to study the influence of the communication bandwidth and latency for MPI applications.

Tuesday, June 24, 2003

4:00pm-5:00pm

P3: Parallel Reservoir Simulation on Intel Xeon HPC Clusters

Kamy Sepehrnoori, University of Texas Austin

Numerical simulation of reservoirs is an integral part of geo-scientific studies, with the goal of optimizing petroleum recovery. In this session, we conduct a series of benchmarks by running a parallel reservoir simulation code on an Intel Xeon Linux cluster and study the scalability while using different interconnects for the cluster. Our results show that the simulator's performance scales linearly from one to 64 single-processor nodes, when using a low-latency, high-bandwidth interconnect. In addition to benchmarking, we describe a process-to-processor mapping approach for dual-processor clusters to improve communication performance as well as overall performance of the simulator.

Wednesday, June 25, 2003

10:30am -11:30am

P4: Geoscience Visualization and Seismic Processing Clusters

Phil Neri, Paradigm

The active development of Linux visualization clusters has led to the notion of associating closely compute-intensive seismic processing

Conference Sessions

and geosciences visualization, notably for the purpose of building and verifying velocity and solid models. The options are to implement cross-system data integration, or to share of a common hardware resource. Practical implementations of the integration model will be presented, based on Paradigm's experience with existing production systems. The use of a CORBA-based distributed data architecture will also be discussed. The common hardware concept, still in the design phase, will be analyzed for its expected benefits, economics and potential problems.

Wednesday, June 25, 2003

1:15pm -12:15pm

P5: Case Study: Deploying Large-scale Seismic Processing Clusters at CGG **Compagnie Generale de Geophysique**

Wednesday, June 25, 2003

4:00pm -5:00pm

P6 Grid Computing In The Energy Industry **Jamie Bernardin, DataSynapse**

Grid computing has attracted significant attention in the current IT environment. What are the business and technical factors driving companies to adopt Grid? In this presentation on Grid computing in Oil & Gas, we will examine, frequently encountered obstacles to deploying a grid computing solution, compare the vision of Grid to the realities of today, identify target deployments for distributed computing solutions in the Oil & Gas sector, and describe the value impact of grid computing. DataSynapse will share case studies from its existing engagements as well as identify specific technical requirements unique to the energy market.

Thursday, June 26, 2003

10:30am-11:30am

P7: Drilling in the Digital Oil Field: High Pay-offs from Linux Clusters

Shawn Fuller, HP

The Oil & Gas industry is required to manage mammoth volumes of complex data for both engineering and scientific requirements in their search for discovering new reservoirs and more cost efficient production methods. Globally deployable high-performance computing systems coupled with best-in-class applications are the keys to success for Oil & Gas companies to excel in their business. This session will cover the areas of technology receiving the most focus: mobility, desktop visualization, scalable and immersive visualization, global collaboration, scalable clustered systems, network storage systems, imaging and printing - covering the full gamut of Oil & Gas IT requirements.

AUTOMOTIVE & AEROSPACE ENGINEERING

Tuesday, June 24, 2003

10:30am-11:30am

E1: Cluster Computing in Space Applications

Eric George, The Aerospace Corporation

This case study will examine how The Aerospace Corporation utilizes cluster computing for a variety of applications in support of high priority national defense programs including the Global Positioning System (GPS) and future missile warning programs. Applications to date have focused on astrodynamics, satellite constellation design, communications network modeling, thermal analysis, and complex scheduling/tasking algorithms. Processing techniques range from Monte Carlo analysis & brute force search operations to genetic algorithms. Research is progressing on implementation of a diverse grid-computing environment at Aerospace.

Tuesday, June 24, 2003

1:15pm -2:15pm

E2: Full Vehicle Dynamic Analysis Using Automated Component Modal Synthesis

Peter Schartz, MSC.Software

Today it is commonplace to attempt to analyze the fully trimmed body of an automobile for its vibration characteristics, over increasing frequency ranges, and on inexpensive computer hardware. The cost effectiveness of RISC based cache processors, combined with upward pressure in the form in large, detailed models, has allowed new software methods to utilize domain decomposition to enable high-level parallelism. A domain decomposition, followed by a component modal synthesis solution, is the bases for Automated Modal Component Synthesis (ACMS) in MSC.Nastran. The solution is described in theory, and its effectiveness is demonstrated by an example taken from today's automotive industry.

Tuesday, June 24, 2003

4:00pm-5:00pm

E3: Using Clusters to Deliver Turn Key CFD Solutions

Paul Bemis, Fluent

While low cost, high performance clusters have been in use since the early 1990's, the application of commercial off-the-shelf CFD software, such as Fluent, to harness these shared nothing architectures has only been viable near the end of that decade. Early implementations required persistent IT department willing to commit the time and

resources necessary to overcome these challenges. Now, however, organizations are able to access a full-featured implementation of Fluent via the Internet in a pay as you go scenario. This session will discuss the problems solved and gains realized by a distributed implementation of Fluent 6.1.

Wednesday, June 25, 2003

10:30am -11:30am

E4: LS-DYNA: CAE Simulation Software on Linux Clusters

Guangye Li, IBM

LS-DYNA is used in a wide variety of simulation applications: automotive crashworthiness & occupant safety; sheet metal forming, military and defense applications, aerospace industry applications, electronic component design. Several years ago, one simulation of a very simplified finite element model needed days to complete on a Symmetric Multiprocessing (SMP) vector computer. With the introduction of Distributed Multiprocessing technology, the MPP (Massively Parallel Processors) version of LS-DYNA can dramatically reduce the turnaround time for the simulation and therefore reduce the time for the automotive design process. We will present the comparison of the scalability of the SMP and MPP versions of LS-DYNA, as well as the comparison of communication networks (Myrinet, Fast Ethernet, Gigabit Ethernet) on Linux clusters.

Tuesday, June 24, 2003

1:15pm -12:15pm

E5: Linux Clusters in the German Automotive Industry

Dr. Karsten Gaier, science + computing AG

After the first German CAE-Linux computer cluster (LCC) was installed in 1999 at DaimlerChrysler for electromagnetic compatibility calculations (EMC), there has been great success in the adoption of LCC. This includes clusters based on 512 CPUs used for crash-calculations running at a major automotive manufacturer. This talk will provide an overview of ways in which Linux clusters are changing the course of CAE in Germany. It will also look at a number of different configurations currently being implemented in some of the world's largest automotive manufacturers.

Wednesday, June 25, 2003

4:00pm -5:00pm

E6: Improving Clusters Through Data Grids in the Automotive and Aerospace Industries

Andrew Grimshaw, Avaki

As the pressure increases to optimize the product design and manufacturing processes

it is critical for the automotive and aerospace industries to give professionals secure access to product and manufacturing information. Data is often located at multiple R&D sites and suppliers, regardless of location. Additionally, product developers require more and more processing power, delivered via clusters that are not effective unless they can provide access to the data they need. This session will examine the most significant data challenges facing today's automotive and aerospace companies and how Grid technology impacts the engineering and manufacturing process.

Thursday, June 26, 2003

10:30am-11:30am

E7: Scrutinizing CFD Performance on Multiple Linux Cluster Architectures

Thomas Hauser, Utah State University

Linux cluster supercomputers are a cost-effective platform for simulating fluid flow in engineering applications. However, obtaining high performance on these clusters is a non-trivial problem, requiring tuning and design modifications to the Computational Fluid Dynamics (CFD) codes. Investigations in optimizing CFD codes on Linux cluster platforms will be presented. Detailed performance results of two CFD codes on a wide range of cluster architectures, including Pentium and Athlon, Intel Itanium and the AMD Opteron, will be analyzed. The single and multi-processor performance of these codes on different cluster architectures will be compared and means of improving performance discussed.

Thursday, June 26, 2003

1:15pm -2:15pm

E8: Managing CAE Simulation Workload in Cluster Environments

Michael M. Humphrey, Altair

Automotive manufacturers are beginning to capitalize on workload management software to get the most out their numerically intense computing environments. Workload management software is middleware technology that sits between your compute-intensive applications - such as ABAQUS, ANSYS, FLUENT, LS-DYNA, NASTRAN and OPTISTRUCT - and your network hardware operating systems. The software schedules and distributes all types of application runs (serial, parallel, distributed memory, parameter studies, big memory, long running, etc.), on all types of hardware (desktops, clusters, supercomputers and even across sites). This presentation will describe the current capabilities of PBS Pro workload management software as a middleware enabler for robust system design.

Travel and Accommodations

Special rates have been negotiated for CWCE 2003 attendees at hotels located near the San Jose Convention Center. All hotel requests should be submitted directly to the San Jose Convention and Visitor's Bureau (CVB) by June 13, 2003. Please make your reservation through our online form.

Fairmont San Jose (Headquarter Hotel)	\$149 + tax (single/double)
San Jose Marriott	\$139 + tax (single/double)
Hilton San Jose	\$129 + tax (single/double)

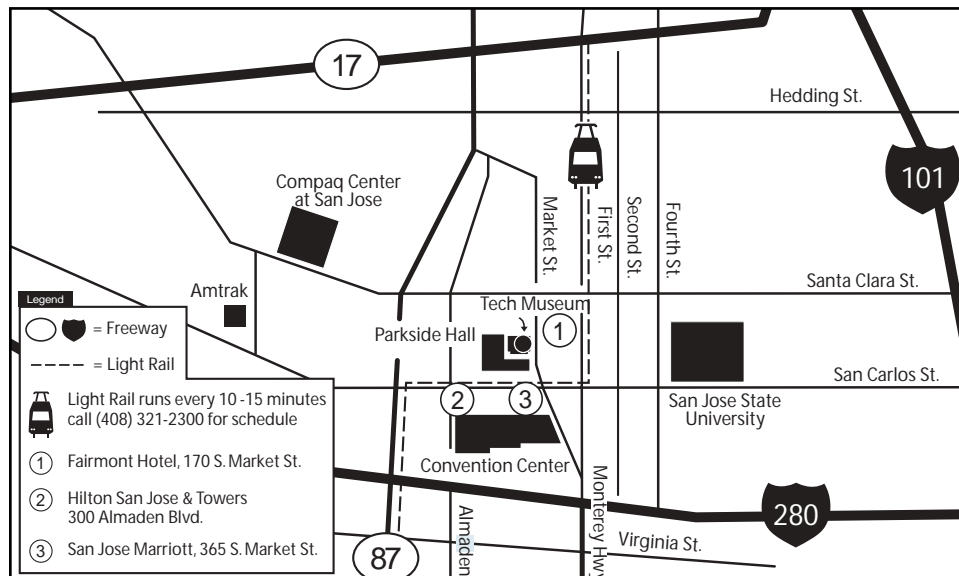
You may contact the CVB directly by calling (408) 792-4168. You may also fax requests to (408) 293-3705. Please do not contact the hotel directly as these special rates are only available through the CVB. Room rates are guaranteed provided that rooms are still available.

HURRY! These discounted rates may not be available after June 13.

Airline Reservations: American Airlines is the official carrier for the ClusterWorld Conference and Expo. To take advantage of the special fares that are being offered to CWCE 2003 attendees, please contact American Airlines directly at (800) 433-1790 and refer to authorization code # A4463AX. Best flights will sell out quickly and seats are limited.

Car Rental: Contact AVIS directly at (800) 331-1600 or online at www.avis.com and reference discount code # J998405 for the best rates available.

Convention Center Location: The San Jose Convention Center is located at 150 W. San Carlos Street, between Almaden Boulevard and Market Street, in San Jose, CA. The Convention Center parking garage entrances are located on Almaden Boulevard and Market Street. In addition, there are more than 21,000 parking spaces in downtown San Jose; parking locations are subject to change due to downtown development. For the most up-to-date information on downtown parking including free parking and validation programs, public transportation and the DASH shuttle, visit www.sjdowntownparking.com.



General Information

CONFERENCE DISCOUNTS

Academia and government employees are eligible to receive a 50% discount off all conference registration packages. To take advantage of this discount, you must either fax or mail the ClusterWorld Conference Registration form to the fax number or address below by June 13, 2003.

Fax the completed form to:
ClusterWorld Conference & Expo
c/o ExpoExchange Registration
(301) 694-5124

Mail the completed form to:
ClusterWorld Conference & Expo
c/o ExpoExchange Registration
PO Box 590
Frederick, MD 21705-0590

A valid ID is required to receive the academia and government employee discount. If you wish to take advantage of the discounted pricing, please be sure to indicate as such on the Conference Registration Form, then attach a copy of your identification credentials and submit the completed form via fax or mail. Valid forms of identification include:

- Educational Institution identification
- Institutional letter verifying educator/student's employment or enrollment status
- Health Insurance identification displaying Educational Institution name
- Government agency identification

Group Discount

Register for four Conference Packages and receive the fifth of equal or lesser value for FREE! Please mail or fax all of your Registration forms together and identify which registrant is FREE. Please note that this special group discount is not available online.

RECEIVING YOUR BADGE

If you are a US attendee and you register on or before May 23, 2003 you will receive your badge in the mail. To receive your badge holder, simply bring your badge with you on-site to any Badge Holder Pick-Up counter, located in the Main Lobby of the San Jose Convention Center. All US attendees that register after May 23, 2003 must pick up their badge on-site during registration hours.

If you are an International attendee, your badge will not be mailed to you. In order for you to receive your badge, please bring your registration confirmation letter with you on-site to the Conference Registration counter, located in the Main Lobby of the San Jose Convention Center.

REGISTRATION AND EVENT POLICY

Registration Cancellation and Substitutions

If you need to cancel, you may do so for a full refund until June 18, 2003. Attendees who register prior to, or after the deadline date who do not cancel in writing by the deadline date are liable and will be charged for the full registration fee. Sorry, no refunds or letters of credit are available after this date. You may fax or mail your cancellation request in order for it to be processed.

Fax the completed form to:
ClusterWorld Conference & Expo
c/o ExpoExchange Registration
(301) 694-5124

Mail the completed form to:
ClusterWorld Conference & Expo
c/o ExpoExchange Registration
PO Box 590
Frederick, MD 21705-0590

Written requests for a downgraded pass must be received no later than June 18, 2003 for a full refund on the difference of registration fees between the value of the original and downgraded pass. Requests received after June 18, 2003 will receive a letter of credit for the 2004 CWCE issued for the difference in pass values. Upgrade pass requests must be submitted in writing and faxed to (301) 694-5124 along with payment information for the difference in value.

Substitutions are allowed only with the written permission of the original registrant. Please mail your substitution request to the above address, or fax to (301) 694-5124.

REGISTER NOW!

3 EASY WAYS TO REGISTER

ONLINE:

June 18, 2003
Visit www.clusterworldexpo.com
(Advanced discount pricing ends
May 23, 2003)

MAIL:

By June 13, 2003
Complete the registration form and send to:
ClusterWorld c/o ExpoExchange Registration
PO Box 590
Frederick, MD 21705-0590

FAX:

By June 13, 2003
Complete the registration form
and fax to: (301) 694-5124
(Faxed registrations must include
credit card information)

REGISTRATION PACKAGES

		Early Bird Discount On or Before May 23, 2003	May 23, 2003 to June 18, 2003	Onsite June 22-26, 2003
Platinum Pass Best Value! (PL)	<ul style="list-style-type: none"> 3 Days of Sessions 2 Half-Day Tutorials (Attendees must register for specific tutorials) Birds-of-a-Feather Sessions Keynotes Feature Presentations Exhibits 	\$895	\$995	\$1,095
3 Day ClusterWorld Conference Pass (3D)	<ul style="list-style-type: none"> 3 Days of Sessions Birds-of-a-Feather Sessions Keynotes Feature Presentations Exhibits 	\$595	\$645	\$695
1 Day ClusterWorld Conference Pass (1D)	<ul style="list-style-type: none"> 1 Day of Sessions Birds-of-a-Feather Sessions Keynotes Feature Presentations Exhibits 	\$225	\$275	\$325
Half-Day Tutorials (TH)	<ul style="list-style-type: none"> Choose up to 2 (Attendees must register for specific tutorials) Keynotes Feature Presentations Exhibits 	\$300 each	\$350 each	\$400 each
Exhibits Pass (EO)	<ul style="list-style-type: none"> Exhibits Keynotes Feature Presentations 	Free	Free	\$50

TUTORIAL KEY

SYSTEMS TRACK TUTORIALS

TS1: Introduction to Linux Clusters - Design
Alternatives, Configuration and Installation

TS2: The Linux Cluster Software Stack - Operations
from the System Administration Perspective

APPLICATIONS TRACK TUTORIALS

TA1: Programming in the Linux Cluster Environment I -
Tools and Program Optimization Techniques

TA2: Programming in the Linux Cluster Environment II -
Parallel Programming with MPI

Register online with your priority code by May 23, 2003, and save!

CONFERENCE JUNE 23-26, 2003

EXPOSITION JUNE 24-26, 2003

SAN JOSE CONVENTION CENTER, SAN JOSE, CA

ClusterWorld™
Conference & Expo

Check here if you require special services. (Attach a written description of your needs.)

ATTENDEE INFORMATION

FIRST NAME	M.I.	LAST NAME
TITLE	COMPANY	
STREET ADDRESS, P.O. BOX, APT#, SUITE, MAIL STOP, ETC.		
CITY	STATE/PROVINCE	
ZIP/POSTAL CODE	COUNTRY	FAX
PHONE	EMAIL	

EMAIL

I would not like to receive product information or news from ClusterWorld Conference and Expo exhibitors or approved third parties via email.

REGISTRATION PACKAGE SELECTION

Please select for the package choices below and the total costs at the bottom. Refer to the opposite page for package descriptions, tutorial codes and pricing. If selecting the Platinum (PL) or Exhibits (EO) pass please do not select any other packages.

Package: Platinum Pass (PL) ClusterWorld 1 Day Pass (1D) ClusterWorld 3 Day Pass (3D) ClusterWorld Half-Day Tutorial Codes (T1) Exhibits Pass (EO)

(Please indicate)

Group Discounts: Register for four conference packages and receive the fifth of equal or lesser value for FREE (all forms must be received together by mail/fax ONLY)

Check here for Government Employee, Academic or Student discount eligibility. Attach a photocopy of your valid identification credentials.

ATTENDEE PROFILE (Must be completed to process your registration)

<p>1) Which category best describes your primary JOB TITLE:</p> <input type="checkbox"/> 1. Corporate Management (CEO, CTO, CFO, President) <input type="checkbox"/> 2. Research & Development Management <input type="checkbox"/> 3. IT Director / Manager <input type="checkbox"/> 4. Lab Director / Manager <input type="checkbox"/> 5. Senior Scientist / Staff Scientist <input type="checkbox"/> 6. Software Engineer / Programmer <input type="checkbox"/> 7. Network Engineer / Manager <input type="checkbox"/> 8. Systems Engineer / Analyst <input type="checkbox"/> 9. Manufacturing / Engineering Staff <input type="checkbox"/> 10. Consultant <input type="checkbox"/> 11. Government Agency Executive <input type="checkbox"/> 12. Academic Head / University Faculty / Professor <input type="checkbox"/> 13. Student <input type="checkbox"/> 14. Other (Please Specify)	<p>2) Which category best describes the BUSINESS or INDUSTRY in which you work?</p> <input type="checkbox"/> 1. Academia / Education <input type="checkbox"/> 2. Aerospace <input type="checkbox"/> 3. Automotive <input type="checkbox"/> 4. Biotechnology <input type="checkbox"/> 5. Digital Content Creation <input type="checkbox"/> 6. Finance / Insurance / Banking <input type="checkbox"/> 7. Government (Defense) <input type="checkbox"/> 8. Government (Energy) <input type="checkbox"/> 9. Government (Other) <input type="checkbox"/> 10. Manufacturing / Design (Computers / Communications) <input type="checkbox"/> 11. Manufacturing / Design (Other) <input type="checkbox"/> 12. Petroleum / Geophysical Exploration <input type="checkbox"/> 13. Pharmaceuticals <input type="checkbox"/> 14. Scientific Visualization <input type="checkbox"/> 15. Other (Please Specify)	<p>3) What is the SIZE of your organization?</p> <input type="checkbox"/> 1. Under 50 <input type="checkbox"/> 2. 50 - 99 <input type="checkbox"/> 3. 100 - 499 <input type="checkbox"/> 4. 500 - 999 <input type="checkbox"/> 5. 1,000 - 4,999 <input type="checkbox"/> 6. 5,000 - 9,999 <input type="checkbox"/> 7. 10,000 or more	<p>4) What is your organization's TOTAL ANNUAL IT BUDGET for all cluster-related hardware, software and services?</p> <input type="checkbox"/> 1. \$0 - \$25,000 <input type="checkbox"/> 2. \$25,001 - \$100,000 <input type="checkbox"/> 3. \$100,001 - \$500,000 <input type="checkbox"/> 4. \$500,001 - \$1,000,000 <input type="checkbox"/> 5. \$1,000,000 - \$5,000,000 <input type="checkbox"/> 6. Over \$5,000,000	<p>5) What is YOUR ROLE in the purchasing process for all cluster-related hardware, software and services?</p> <input type="checkbox"/> 1. Evaluate or recommend products <input type="checkbox"/> 2. Specify vendors <input type="checkbox"/> 3. Approve purchases <input type="checkbox"/> 4. No role	<p>6) Do you plan to evaluate products at this conference for purchase within the next year?</p> <input type="checkbox"/> 1. Yes <input type="checkbox"/> 2. No <input type="checkbox"/> 3. Maybe	<p>7) Which PRODUCTS OR SERVICES are you interested in purchasing? (check all that apply)</p> <input type="checkbox"/> 1. Bioinformatics Software <input type="checkbox"/> 2. CAD / CAE / CFD Software <input type="checkbox"/> 3. Database Software <input type="checkbox"/> 4. Digital Content Creation / Rendering Software <input type="checkbox"/> 5. Distributed / Grid Computing Software <input type="checkbox"/> 6. High-Speed Network Interconnects <input type="checkbox"/> 7. IA64 Based Systems <input type="checkbox"/> 8. X86-64 Based Systems <input type="checkbox"/> 9. Non-IA Based 64-Bit Systems <input type="checkbox"/> 10. IA32 Based Systems <input type="checkbox"/> 11. Imaging Software <input type="checkbox"/> 12. Network Infrastructure (Routers / Switches) <input type="checkbox"/> 13. Network / Systems Management (Hardware) <input type="checkbox"/> 14. Network / Systems Management (Software) <input type="checkbox"/> 15. Security Software <input type="checkbox"/> 16. Simulation Software <input type="checkbox"/> 17. Storage Products (Hardware) <input type="checkbox"/> 18. Storage Products (Software) <input type="checkbox"/> 19. Visualization Tools <input type="checkbox"/> 20. Parallel Programming Compilers <input type="checkbox"/> 21. Parallel Programming Profilers / Debuggers <input type="checkbox"/> 22. Other (Please Specify)	<p>8) HOW MANY SERVERS make up your organization's cluster?</p> <input type="checkbox"/> 1. Under 50 <input type="checkbox"/> 2. 50 - 99 <input type="checkbox"/> 3. 100 - 199 <input type="checkbox"/> 4. 200 - 499 <input type="checkbox"/> 5. 500 - 999 <input type="checkbox"/> 6. 1000 - 1999 <input type="checkbox"/> 7. 2000 + <p>9) Would you like to receive a complimentary subscription to ClusterWorld Magazine?</p> <input type="checkbox"/> 1. Yes <input type="checkbox"/> 2. No <p>In lieu of a signature, audit requirements dictate that we must ask for a personal identifier. Please indicate your eye color _____</p>	<p>10) Please indicate your housing plans?</p> <input type="checkbox"/> 1. Use my own travel agent <input type="checkbox"/> 2. Book accommodations myself <input type="checkbox"/> 3. Use ClusterWorld Conference & Expo Housing <input type="checkbox"/> 4. Live Locally <input type="checkbox"/> 5. Staying with relatives / friends <input type="checkbox"/> 6. Don't know yet
--	--	---	--	---	--	---	---	---

PAYMENT INFORMATION (Payment must accompany form for registration to be completed)

Conditions: All Registration fees and credentials are non-transferable. Discounts on registration fees are valid on NEW registrations ONLY and must be redeemed/noted at time of registration. No refunds or credits will be issued for a discount after the initial registration. A copy of the order must be provided at the time of registration. A \$20.00 fee will be charged for all returned checks.

TOTAL Amount \$ _____ Check Enclosed (Make payable to ClusterWorld Conference & Expo and enclose the registration form in the envelope)

MasterCard VISA American Express

ACCOUNT NUMBER	EXPIRATION DATE
FIRST NAME	M.I. LAST NAME

PRIORITY CODE: Enter your priority code from promotional mailing or email. If you do not have a code, please enter None.

Cancellation and Substitutions: If you need to cancel, you may do so for a full refund until June 18, 2003. Attendees who register prior to, or after the deadline date, who do not cancel in writing by the deadline date are liable and will be charged for the full registration fee. Sorry, no refunds or letter of credit are available after this date. Please fax your cancellation request to 301-694-5124 or mail your request to ClusterWorld Conference & Expo, c/o ExpoExchange Registration, PO Box 590, Frederick, MD 21705-0590. Written requests for a downgraded pass must be received no later than June 18, 2003 for a full refund on the difference of registration fees between the value of the original and downgraded pass. Requests received after June 18, 2003 will receive a letter of credit for the 2004 CWCE issued for the difference in pass values. Upgrade pass requests must be submitted in writing and faxed to 301-694-5124 along with payment information for the difference in value. Substitutions are allowed only with the written permission of the original registrant. Please mail your substitution request to the above address, or fax to 301-694-5124. Sorry, no one under the age of 18 is allowed on the Expo floor.