



Axceleon Enfuzion for LC Series

Session 3279



Dan Cox
Manager – ISS HPTC Programs
Hewlett-Packard



Axceleon Enfuzion for LC Series – Topics




- * HP HPTC Portfolio and LC Series Partner Driven Software Solutions
- * Enfuzion Overview
- * Q&A's

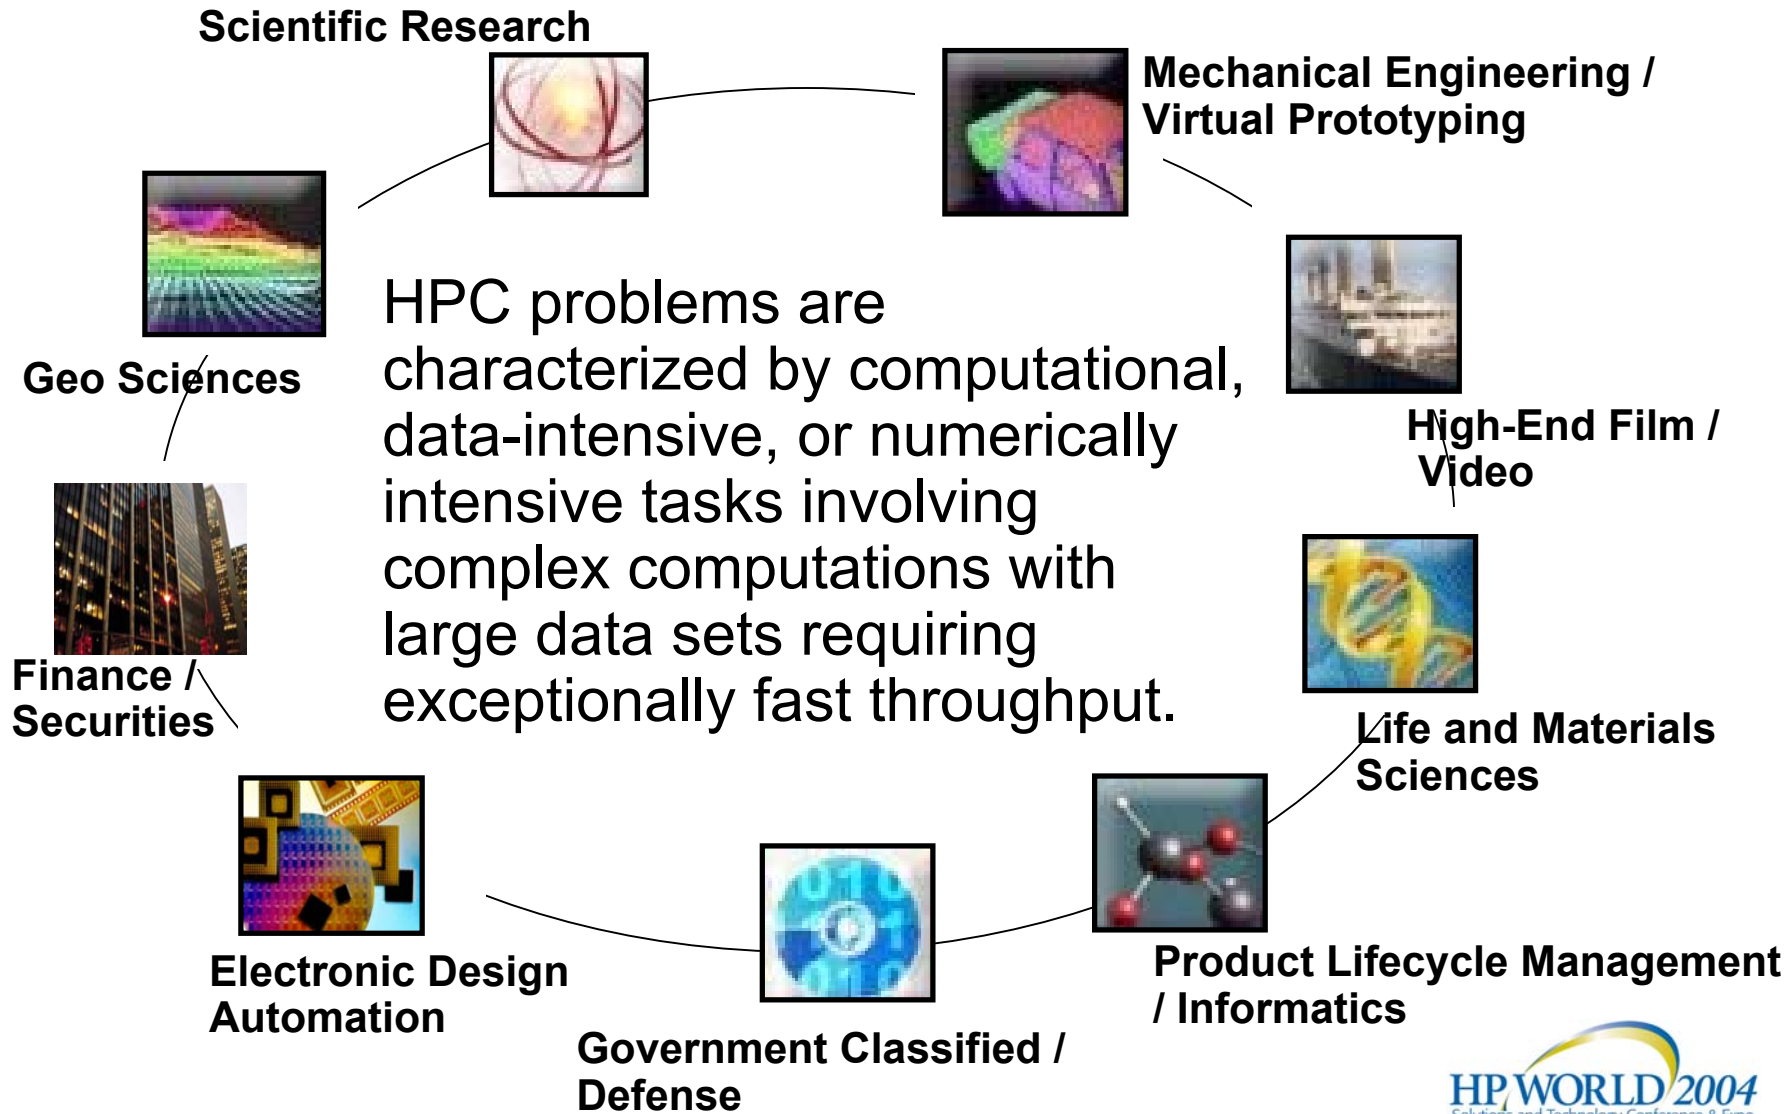
HPTC and new technologies at HP

“High performance technical computing is a strategic focus for HP. It is here that new computing paradigms are created and the applications they enable become early indicators of the general commercial applications that will follow.”

Carly Fiorina
Chairman and CEO, HP



HPC @ HP



HP's HPC Platforms



LC Series 1000/2000
HPTC clusters



LC Series 3000
HPTC cluster



XC3000
HPTC cluster



XC6000
HPTC cluster



HP Integrity
rx1600



HP Integrity
rx2600



HP Integrity
rx4640



HP Integrity
rx7620



HP Integrity
rx8620



HP Integrity
Superdome



HP ProLiant
DL140/DL145




HP ProLiant
DL360/DL380



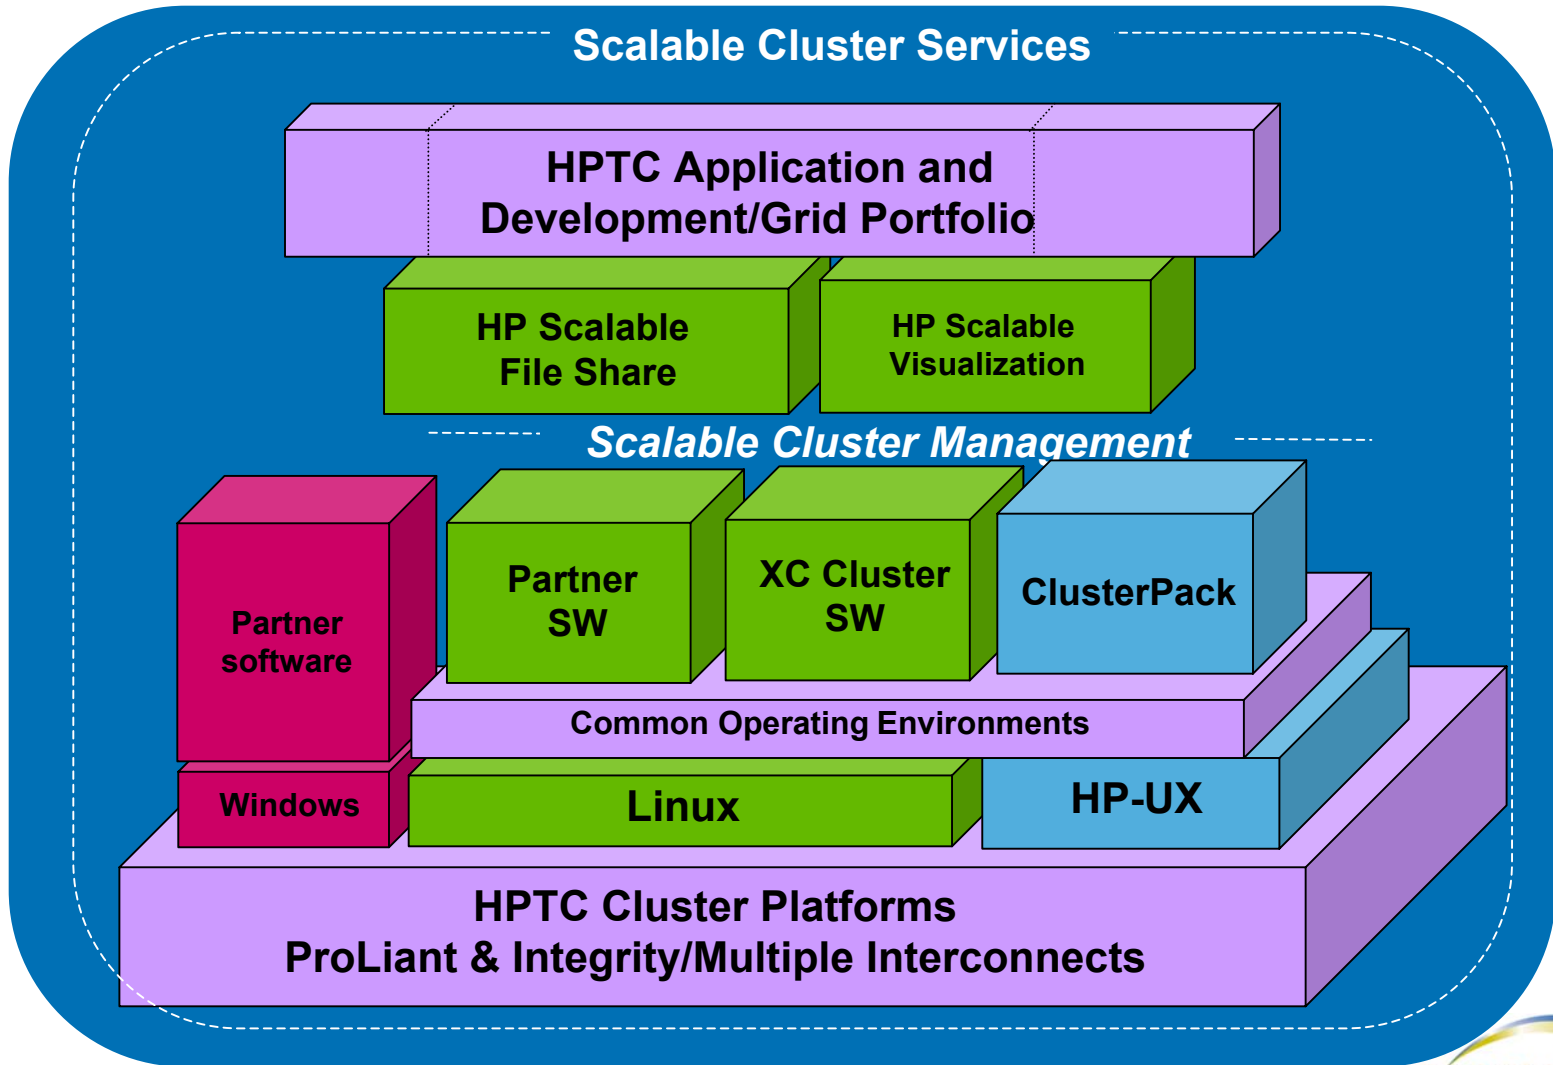
HP ProLiant
BL20p/BL30p



HP ProLiant
DL585



HP's HPC Software Architecture Portfolio



HP's HPC solution strategy –

- Goal:

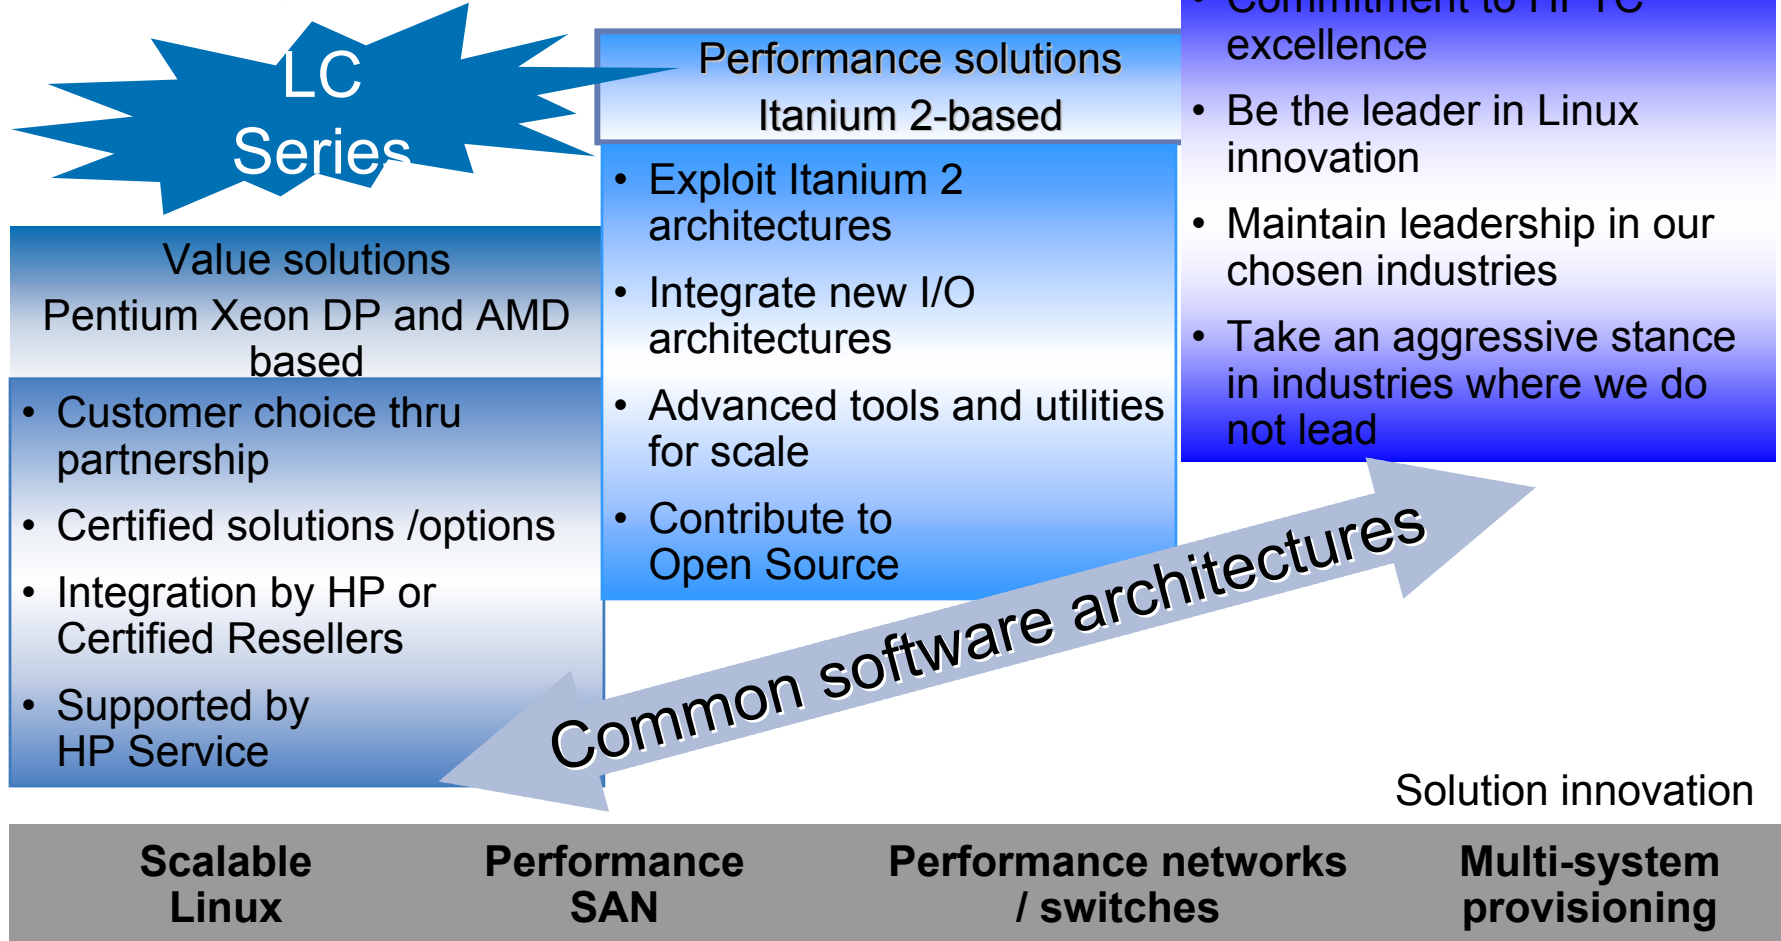
- Deliver complete turn-key integrated solutions that are designed specifically for customer needs

- Challenge:

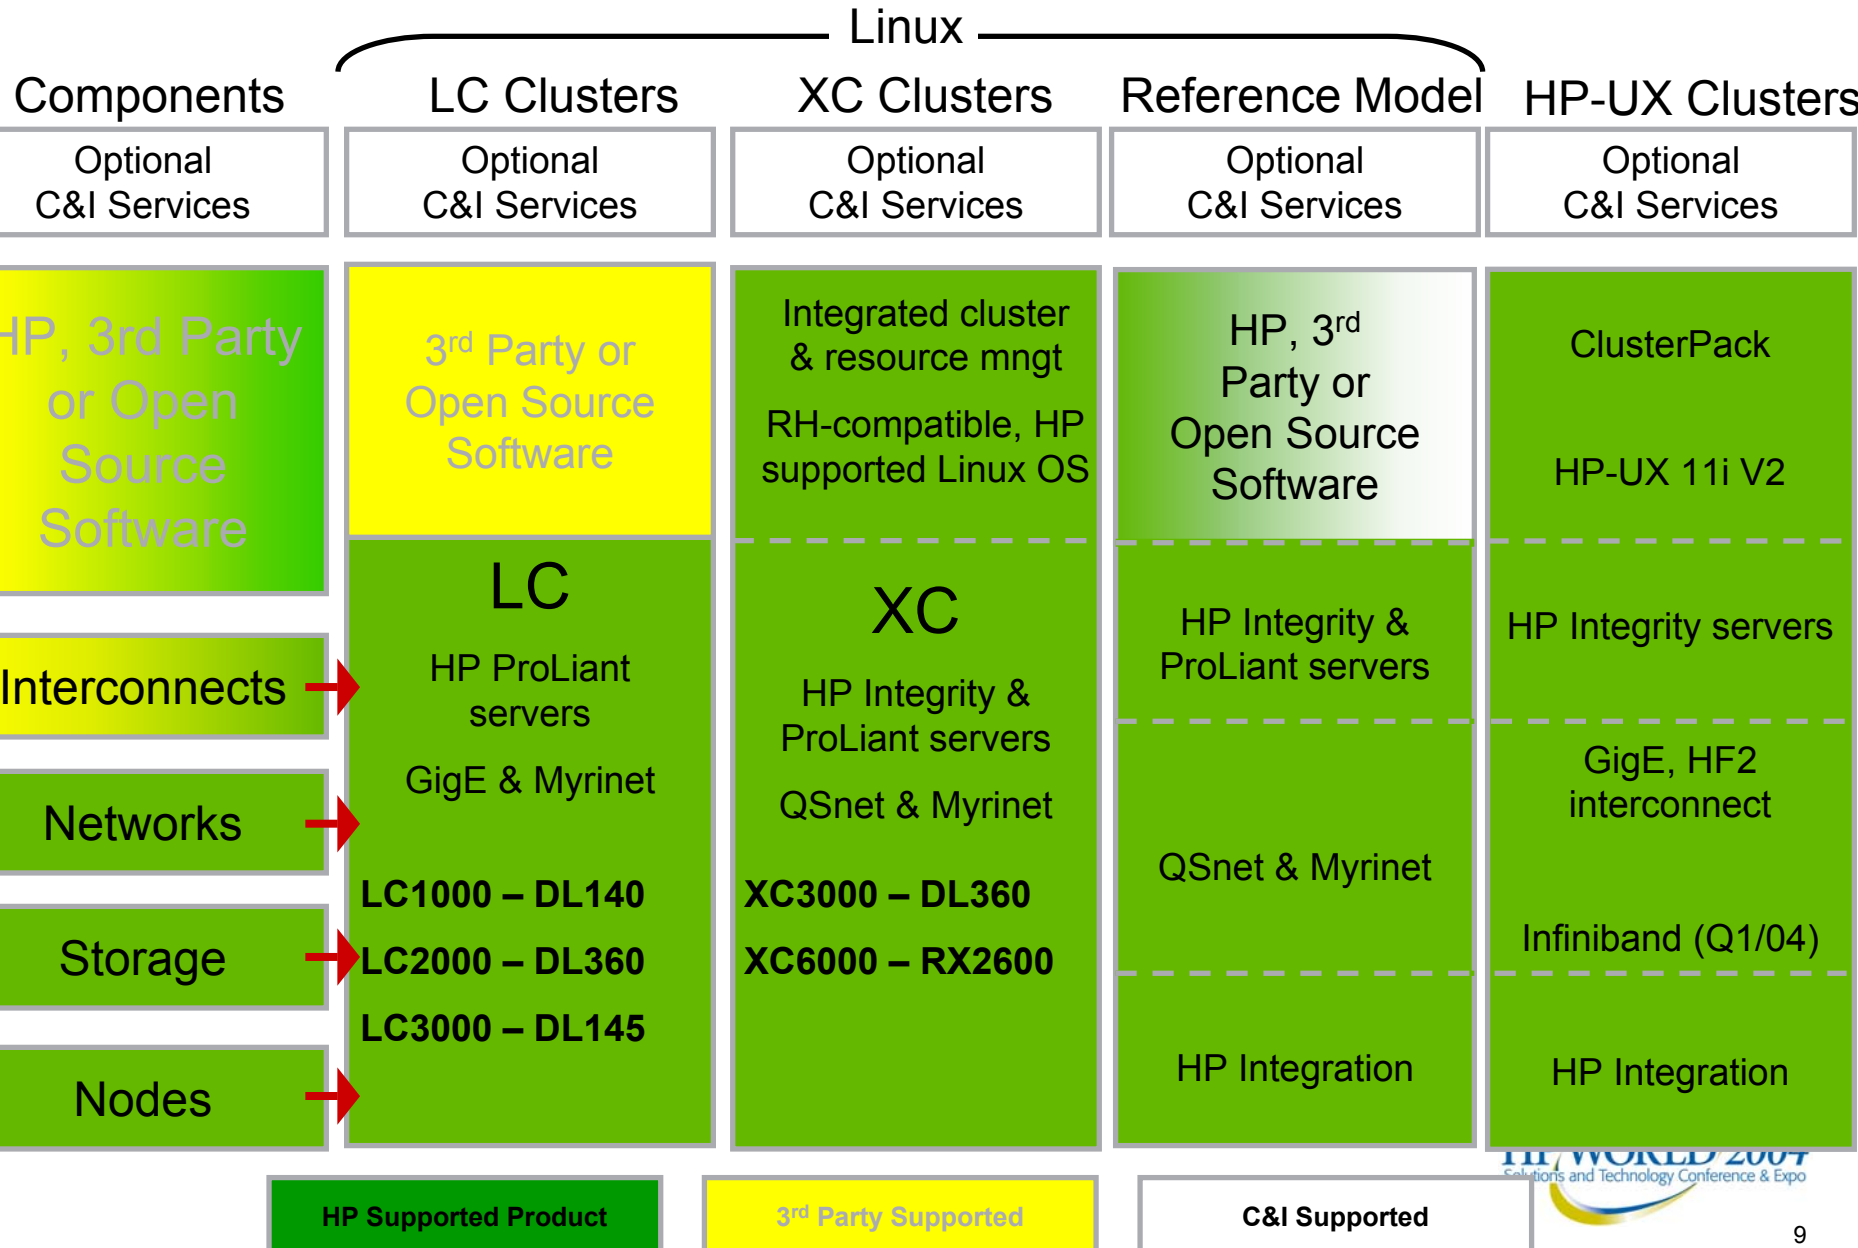
- The key is tight linkages with application partners and a process for design, build, delivery, and support
 - Application codes are complex to integrate as hardware components, maybe more so!
 - Software partners and hardware companies have not been taking ownership for this!
- A viable production cluster solution requires a unique set of skills and products only available through a partnership with leaders in their respective fields and Professional Services trained and aware of HP solutions and alternatives

HP's HPC product delivery

Deliver unprecedented capabilities through pervasive technologies and partnerships



HP's High Performance Cluster Choices



LC Series Overview

- LC for 'loosely coupled'
- Three new orderable 32-node SKU solutions
 - **LC1000 Series** — ProLiant DL380 Control node and 32 ProLiant DL140 Compute nodes
 - **LC2000 Series** — ProLiant DL380 Control node and 32 ProLiant DL360 Compute nodes
 - **LC3000 Series** — ProLiant DL380 Control node and 32 ProLiant DL145 Compute nodes
- Common Packaging for ease of ordering, manufacturing and support
 - Out of Band unique due each platform device, IMPI 1.0, 1.5, and iLO



- Optional software from HP and ISV partners
 - Linux and Windows 2003 HPC
 - Checkpoint Restart
 - HPTC Cluster and GRID Mgmt
 - Development ToolKits
 - Job Management
- Optional Storage Bundles for NFS and GFS Storage support
 - Red hat Sistine
 - PolyServe Matrix Server
- Factory integrated and tested in Houston and Erskine
- Complete suite of service and services offerings

HP LC Series Family Overview


LC Series Design and Configuration Guide

Reference Models for flexibility:

- ✓ Any server model (2.8, 3.06, 3.2...)
- ✓ Any compute node count up to 128
- ✓ More than one control node
- ✓ Storage options
- ✓ Change switch types/size
- ✓ Modified cabling
- ✓ Customized software installs
- ✓ Component location changes

Reference Models require:

- ✓ Purchase of the Configuration Resource Kit
- ✓ Factory presets of servers and switches
- ✓ Cabling consistent with the LC cluster specification
- ✓ Use of components certified for use in LC clusters



Reference Model Benefits:

- ❖ Speed flexibility
- ❖ 8 to 128 node reference design
- ❖ Supports HA Control Node
- ❖ NFS to 48TB GFS options
- ❖ Flexibility and choice
- ❖ Meet specific site need
- ❖ Add infrastructure products
- ❖ Meet specific site need

Reference Model benefits:

- ❖ Fully documented configuration for ease of growth or service
- ❖ Geared to application performance
- ❖ Ease of growth for scale out and serviceability
- ❖ Ease of serviceability and scale out

**Over 180 Tested
Configuration Options**

LC Series Partner Driven Software Portfolio



Applications	Bioscience, Engineering, GeoScience, etc
GRID Mgmt	Axceleon , United Devices
HPTC	Scali
Cluster Reliability	Meiosys– Checkpoint Restart
Job Management	Altair
Development	Engineered Intelligence
Data Mgmt	Red Hat Sistina, PolyServe
Operating System	Red Hat / SUSE Linux, Windows



Certified Partner Based Solutions - Linux



Platform™

- Multiple Linux cluster management choices allow customers to find the 'best fit' based on

- experience with Linux and open source
- on-site capabilities
- ISV application focus
- ISV support for target hardware



- Service programs in place, or in roll-out, for these partners
- Available as integrated solutions from each region



Certified Partner Based Solutions - Windows

axceleon

Platform™

- Microsoft support is beginning thru Microsoft and limited partner approaches
- Multiple application choices can be adapted to Windows 2003 Server High Performance Computing Edition
 - Special Windows 2003 HPC pricing for Compute Nodes
- Partner solutions based on end user needs that HP is partnering with
 - Utilities
 - Bio Life and Materials Science
 - Financial Services Modeling

 **POLYSERVE**



“ HP is the unchallenged leader in delivering HPC solutions and provides the most competitive breadth of Linux and Windows cluster solutions offered. LC Series clusters are available to meet just about any customers' needs. They are configured, tested, and shipped ready to run out of the box. No special ordering process is required, regardless of the node architecture.”

Robert de Sautel
Principal, Harvard Research Group

Harvard Research Group, Inc.
P.O. Box 297
Harvard, MA 01451 - 0297 USA
Tel. (978) 263-3399



Axcelon Enfuzion Overview

Customer Needs

- Need for real time analysis (financial)
- Need to shorten time to market (manufacturing)
- Time pressure, deliver fast (digital media)
- Need to have high work quality (energy)
- Need for parametric studies (government, R&D)
- Need for enormous managed compute power (life sciences)
- Increased/optimized resource use (all)

Customer Business Challenges


- Most job scheduling (JS) solutions do not work well in Heterogeneous environments.
 - Tuned usually to single System/Cluster
- Current JS solutions can be expensive
- Current JS solutions can be hard to deploy and maintain, concerns about cost
- Current JS solutions can lack features and are not optimized for parametric execution
- Ease of use very, very, very important
- Need for automation of “run-submission” process

What is EnFuzion?

- An Enterprise Resource Manager™ that combines HP Clusters and Servers and creates an affordable supercomputer, allowing users to do more
- Simplifies job generation, eases the task of distributing applications across many servers
- Enables faster execution through distribution of jobs
- Reduces complexity of distributed computing involving many machines in heterogeneous environments
- Maximizes utilization of compute resources with minimum user involvement
- Manages files, machines and network, freeing users to do problem solving

EnFuzion History

- Nov'93 – The first Linux Beowulf Cluster?
 - Griffith University, Brisbane, Australia
- Mature, field-proven software since 1996
 - 1996 - Deployed in life sciences
 - 1997 - First commercial release
 - 1998 - Deployed in engineering analysis
 - 1999 - First technology of this type on Wall Street
 - 2002 - Acquired by Axceleon
 - 2003 - EnFuzion 8.0 released



Courtesy of David Abramson, Monash University

Why EnFuzion on LC Series

- Clusters are a growing market in the Enterprise
- Each Cluster needs management tools
- EnFuzion does Enterprise Resource Management
- Why EnFuzion
 - EnFuzion makes your customer's business & technical applications high performance
 - Increase compute power, reduce IT costs
 - Simple to install, easy to use
 - Grid enables applications

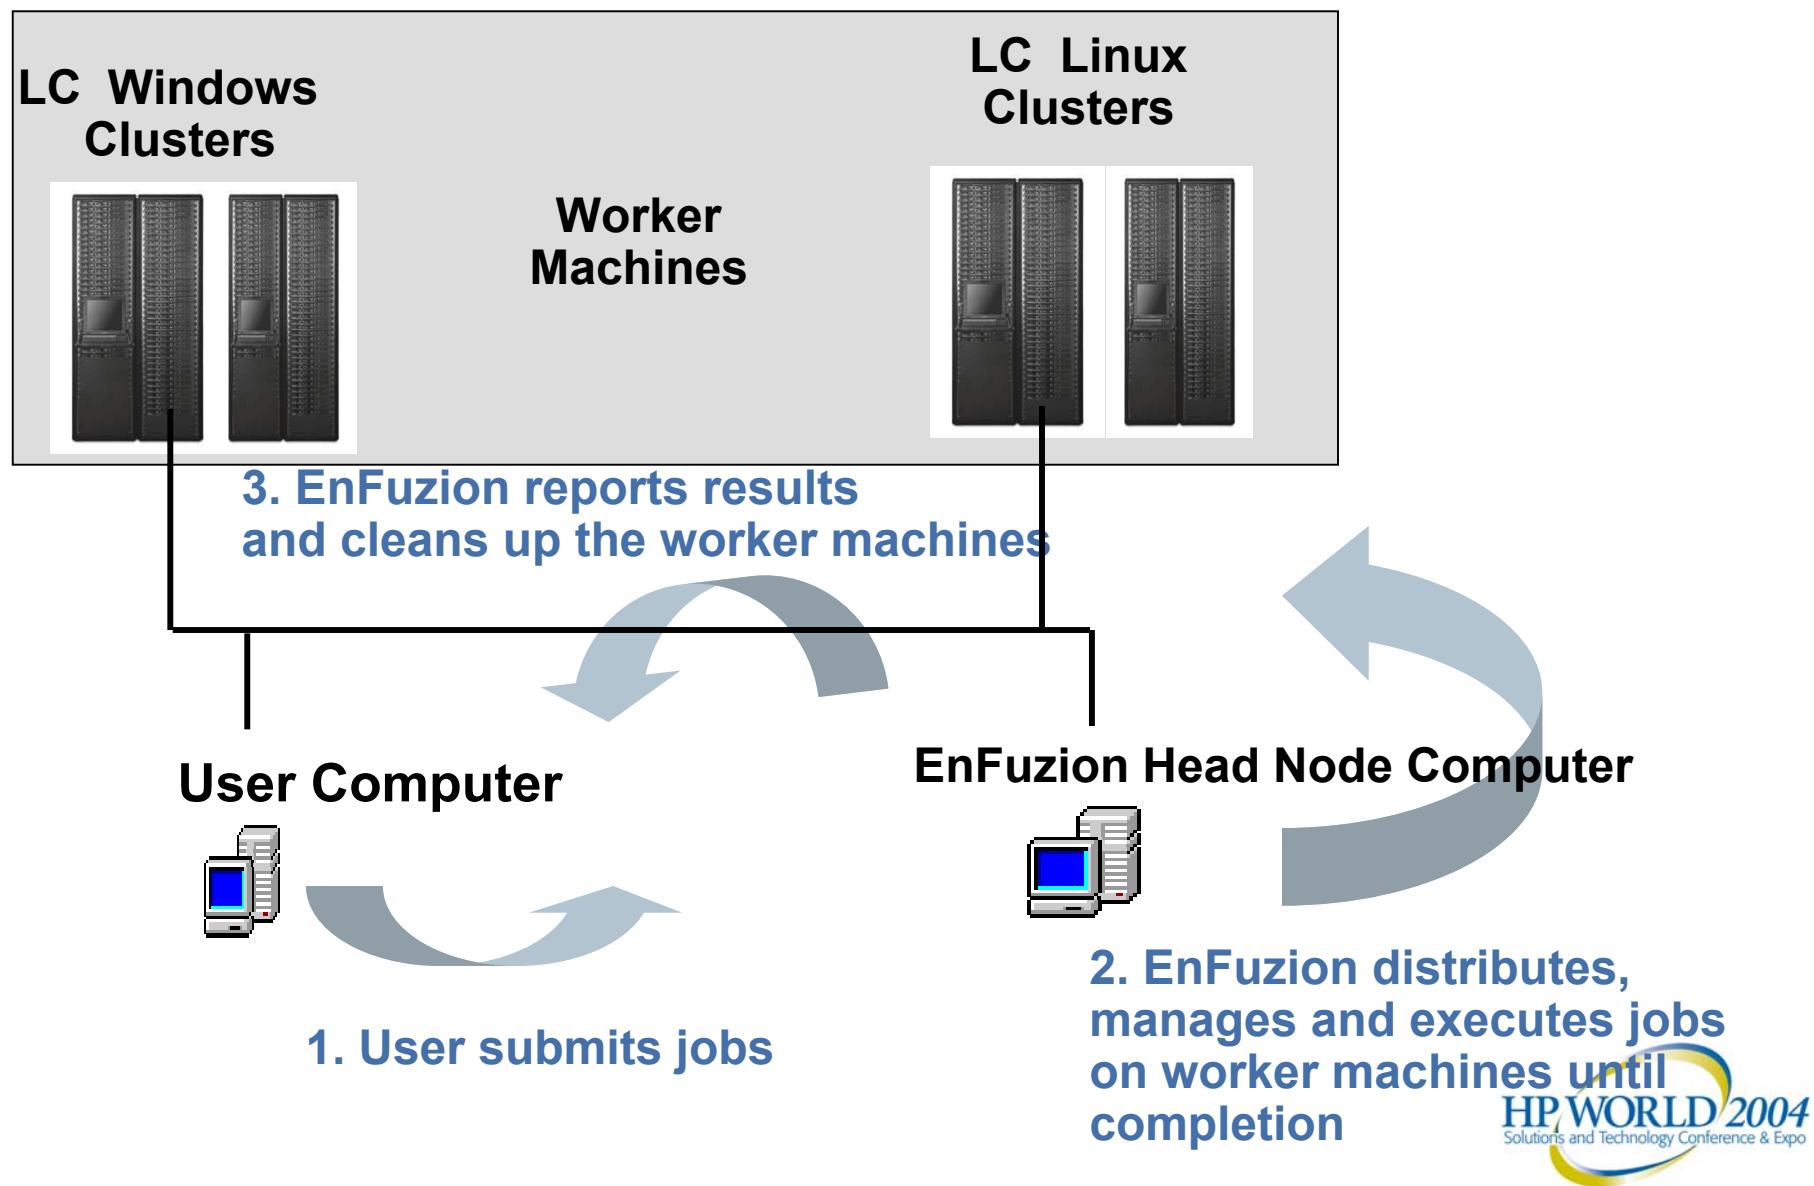
Why EnFuzion the Product

- Optimized for Parametric Execution
 - Same algorithm large number of different inputs or scenarios
- High throughput with low latency
 - Large number of jobs in a very short time
 - Equivalent to a “real time” response
- Application Specific user interface
 - Easy to use GUI
 - Full API
- Works in heterogeneous operating environments, including; IRIX, Unix, Linux, Windows 2000, NT, and XP
- No need to modify existing applications
- Speeds up job's / simulations / rendering up to 1000 times depending on grid size
- Harness the power of every CPU in existing racks/networks
- Easy to deploy, ease to use



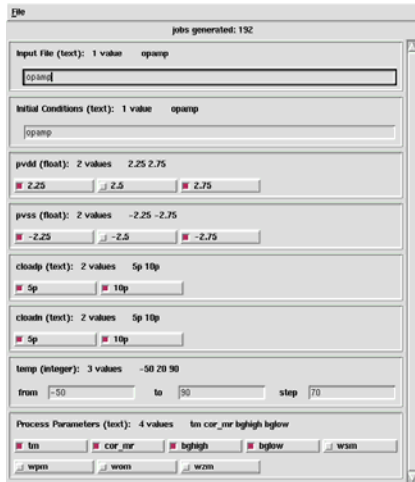
Core Technology

- EnFuzion is software for computational grids
- EnFuzion allows applications to use the combined power of multiple computers.
 - based on parametric job execution, same application, numerous sets of different parameters and scenarios
 - very high performance up to 1,000,000 jobs/minute
 - real time capability down to 0.01s response time
 - easy to deploy and manage, scalable
 - Web based user interface/GUI
 - extensive management and monitoring capabilities
 - automatic load balancing, resource sharing and fault tolerance
- EnFuzion has a wide range of options to integrate with applications from no programming to complete API

How does EnFuzion work?



EnFuzion: the Computational Workbench



K⁺ (Rabi Frequency, Laser Detuning)

Rubidium 1-5/2 isotope : Zero Doppler width

Solutions and Technology Conference & Expo

**Description of
Parameters and
Commands**

“Deploying EnFuzion allowed us to harness the power of hundreds of powerful desktop workstations and servers during times when they would otherwise be sitting idle..“

Michael Liberman,
head of Global Swaps and Derivatives Technology
J.P. Morgan.

Cost savings – more than \$7 million a year at project launch



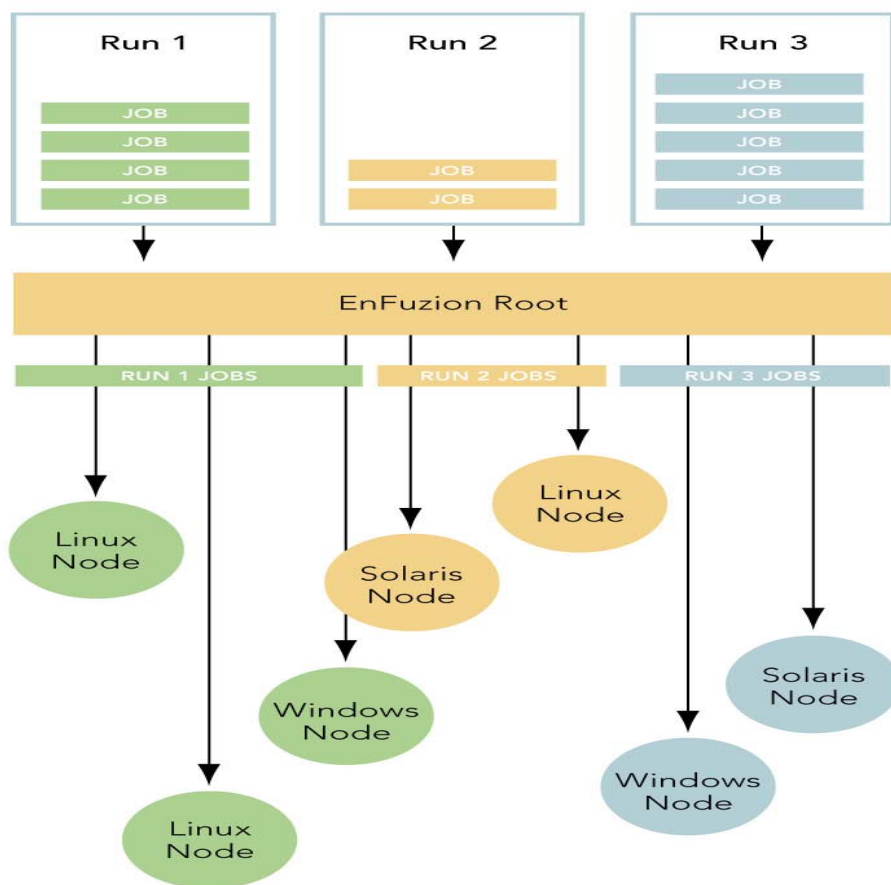
Install and Configure

- **Installation – No Special Requirements**
 - Hardware/software pre-requisites – none
 - Fully self contained, no other applications required
 - Low footprint on the root and node machines
- **Flexible Configuration Options**
 - Can be adapted to any IT infrastructure policies, no need to do anything special to get EnFuzion running
 - Dynamic EnFuzion reconfiguration w/o bringing the system down
- **Heterogeneous Environment Supported**
 - Windows, Linux, all major Unix HP-UX, Tru64
 - 64-bit platforms - Itanium, AMD64, Alpha
 - Can mix and match different platforms

Distributed Resource Management - I

- **Job Management**
 - Automated job scheduling and execution management
 - Transparent restart of jobs from failed machines
 - Differentiated handling of system and job errors
- **File Management**
 - Creation of separated job work areas on worker machines
 - Automated copying of input files to worker machines and output results to the root or the submit computer
 - Automatic clean-up of worker machines
- **Workload Management**
 - Support for multiple users
 - Prioritizing jobs from multiple users/applications according to their priorities
 - User access permissions

Job Management by EnFuzion



- Parametric execution
- Real-time processing 0.01 sec
- Million jobs a minute
- Multi user support with priority
- Easy to use & deploy, user GUI
- Resource management
- Resource monitoring
- Automatic load balancing
- Full API

Intuitive Web Based Dashboard

The screenshot displays the EnFuzion 7.2 web dashboard, which is a web-based interface for managing a cluster. The dashboard is divided into several sections:

- Cluster Information:** Shows the cluster name (host3:48037) and its status (Running). It also displays the uptime (00:12:03:10) and the number of nodes (10).
- Nodes:** A table showing the status of each node in the cluster. The table has columns for Active, Executing, Idle, Busy, and Down. All nodes are currently Active.
- Runs:** A table showing the status of each run. The table has columns for Created, Started, Stopped, Done, and Failed. All runs are currently Done.
- Cluster Log:** A table showing the messages from the cluster. The table has columns for Time, Object, and Message. The messages show the cluster's startup sequence.
- Run List:** A table showing the details of each run. The table has columns for ID, Name, Status, Uptime, Finish In, Priority, Priority Level, Allocated Nodes, Jobs Waiting, Jobs Executing, Jobs Done, Jobs Failed, Job Length, and Total Time. The runs are currently in the 'Started' state.
- Node List:** A table showing the details of each node. The table has columns for Name, Host, Status, Uptime, Executing, Idle, Busy, Downtime, Job Limit, Jobs Executing, Jobs Done, and Job Length. The nodes are currently in the 'Executing' state.

The dashboard is designed to be intuitive and easy to use, with a clear layout and a variety of filters and options for viewing the data.

Distributed Resource Management - II

- **Resource Management**
 - Matching of job requirements with node properties:
 - Application releases
 - OS platforms
 - Application availability
- **Node Management**
 - Restart of EnFuzion software on recovered nodes
 - Dynamic addition and removal of nodes
- **Load Monitoring**
 - Peaceful co-existence with other computer uses, non-EnFuzion applications
 - Configurable to maximize use of dedicated machines or minimize impact to shared machines
 - Use idle machine cycles
- **Application Management**
 - Installation of required applications on remote nodes:
 - Use the latest application version

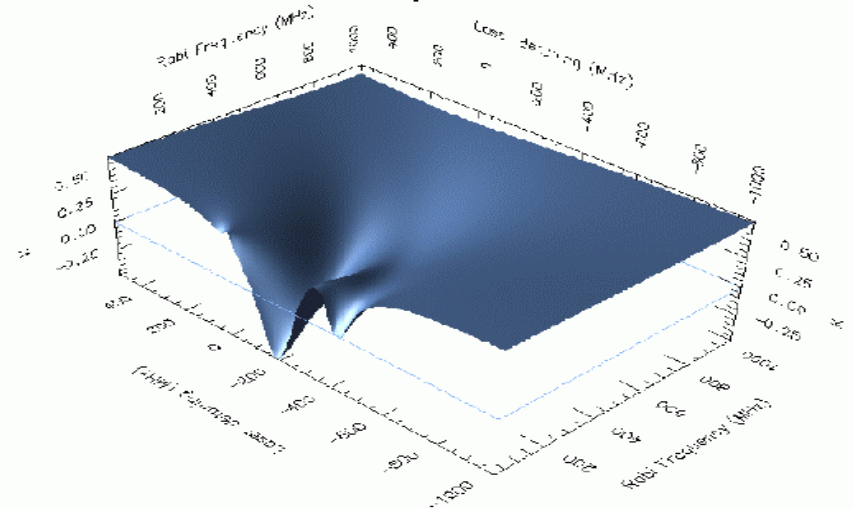
Easy to Manage, Fully Secure

- **Extensive Accounting/Reporting Services**
 - Run and node (job and worker machine) specific reports
 - Hourly, daily, monthly reports
 - Report formats - text, html, cvs or customized
- **Flexibility without Compromise to Security**
 - Extended OS-provided security
 - Network based security
 - Host based security
- **Comprehensive Logging and Debugging Functionality:**
 - Logging of all major events on the root and nodes
 - Additional diagnostic logs and tools
 - Saving application error environment for inspection and analysis

Customer Benefits

- Increased Productivity Leads to Increased Revenue, Top Line ROI

- Faster time to results
 - Reduce time to market
 - Increase quality
- Automate repetitive steps
 - Eliminate manual errors
 - Focus on innovation



- Increased Resource Utilization Leads to Reduced IT Capital and Operational Expenditures, Bottom Line ROI

- Unite computing resources, share machines and applications
 - Maximize use of existing hardware
 - Optimize use of existing software licenses
- Reduce tedious IT maintenance tasks
 - Improve IT productivity
 - Focus IT on business applications

Courtesy of Andrew Lewis, Queensland Parallel Supercomputing Facility

**EnFuzion helped to discover
new science in laser physics**

- Easy to Set up and Run

“Suddenly there was an enormous growth spurt in the amount of work we were able to perform. (EnFuzion) made it possible to multiply the amount of work done by a factor of almost 100.”

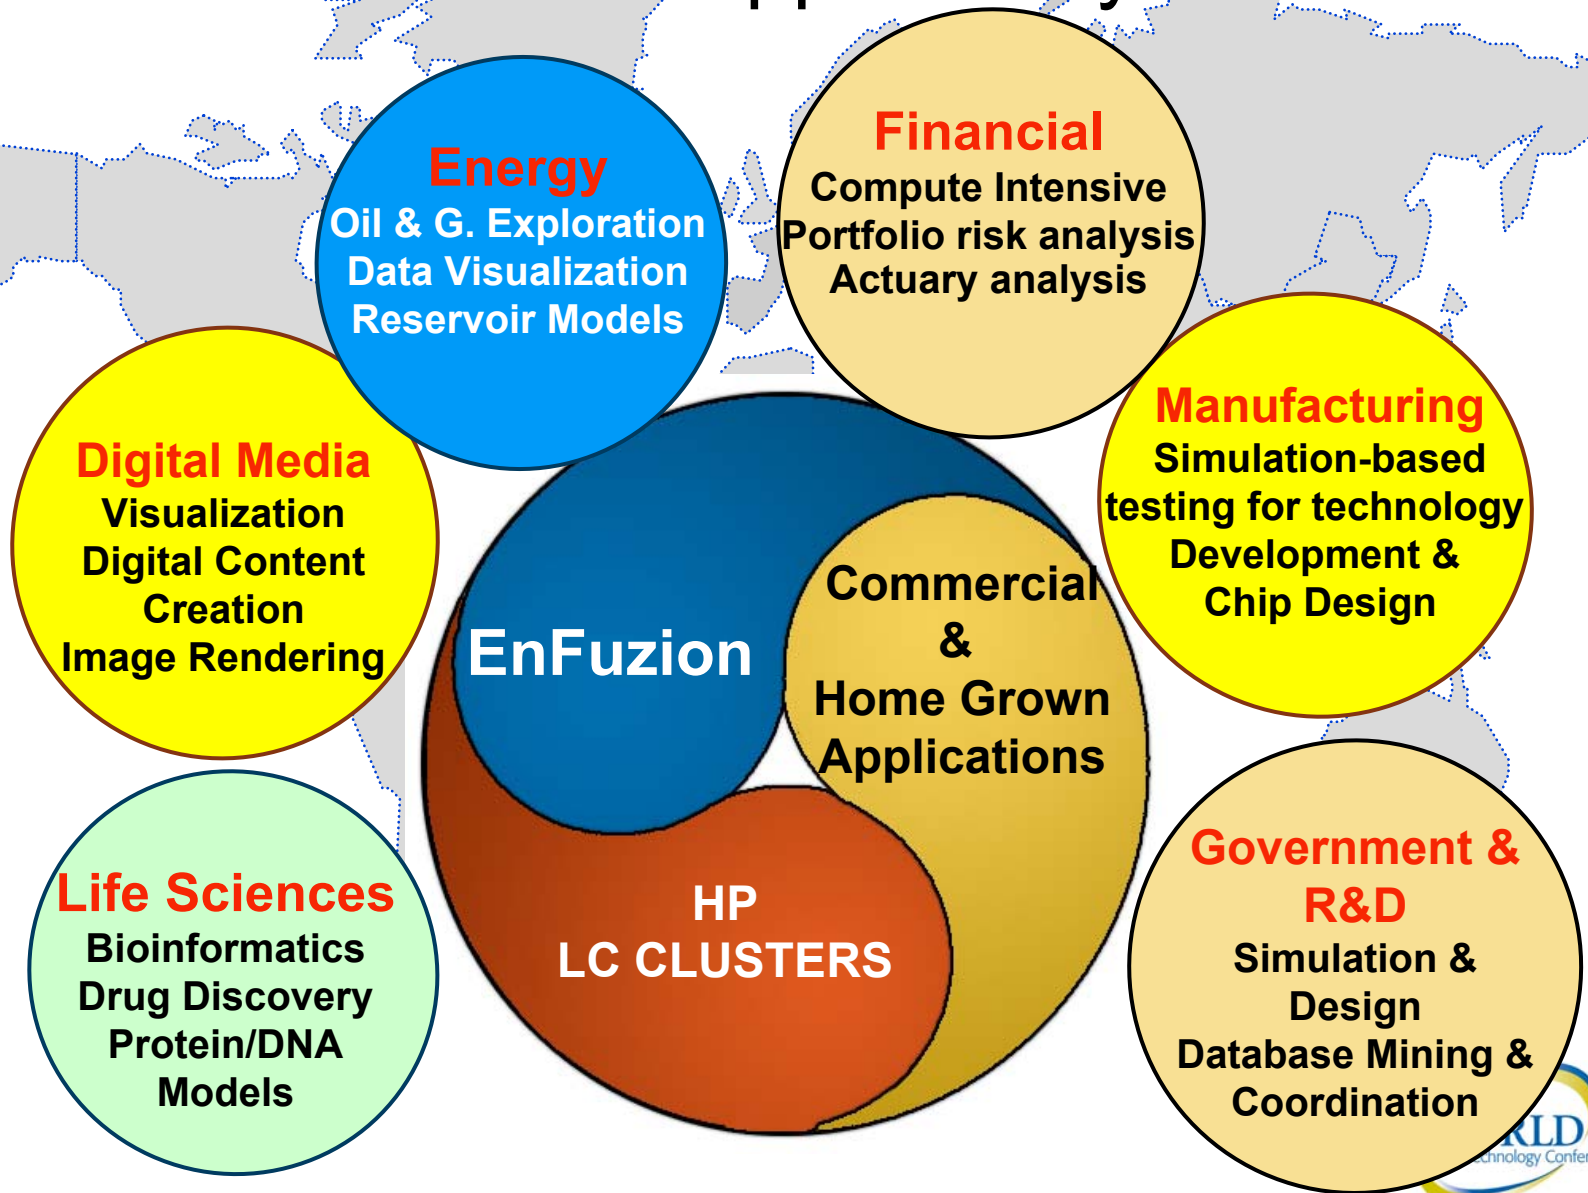
Dr. Alen Varsek

AMP

Australia's largest insurance firm and one of that country's leading banks.



EnFuzion's Applicability



Case Studies

- **Ceres Inc.** - Genomic Research & Technology company
 - Use computational biology to provide accurate solutions for many biological problems
 - Use Blast, psi-blast, PSORT, and SignalP for analysis and mining of predicted three-dimensional protein structures
 - Need to process vast amounts of genomic information
 - Needed a cluster management solution
- **Solution**
 - Use EnFuzion to manage a 50 Node LC Series Cluster
- **Benefits to customer**
 - Automated task management and processing
 - Maximize hardware utilization, minimize IT expenditures
 - EnFuzion **integrated easily** into their environment



Case Studies

- **Powerlink Queensland** - a leading Energy Transmission Co.
 - required increased compute power
 - needed solution for commercial power systems simulation software application



- **Solution**
 - EnFuzion enabled LC Series Cluster deployment
- **Benefits to customer**
 - Reduced simulation time, automated task processing,
 - Turbo-Charged their Electrical Grid Analysis
 - EnFuzion allowed fast, seamless integration with their commercial and home grown software applications

- Contingency analysis in PSS/E normally takes 4 hours to complete – with EnFuzion takes 6 minutes, speedup = 40 times
- Transient stability analysis with different fault durations and scenarios take 7 hours to complete – with EnFuzion takes 46 minutes, speedup = 9.2 times.

“Powerlink’s use of EnFuzion is a great example of a growing trend to consolidate IT costs by utilizing grid computing to manage and execute compute intensive jobs faster and more accurately.”

Bill Claybrook,
Research Director - Aberdeen Group



Q&A

HP delivering choice in HPC

Scale-up, Scale-out, Scale-simply



ProLiant Servers
with 64-bit extensions



Integrity
Servers



- Price/performance leadership with 32/64-bit co-existence
- Highest clock speed, peak performance
- Extensive 32-bit, and emerging 64-bit ecosystems
- Scale-out for simple, highly parallel workloads (2p nodes)
- Linux & Windows

- Price/performance leadership with 32/64-bit co-existence
- 32-bit throughput performance leadership
- Highest bandwidth for sustained performance
- Extensive 32-bit, and emerging 64-bit ecosystems
- Scale-out for moderate workloads (2p/4p nodes)
- Linux & Windows

- Highest performance 64-bit processor core for sustained performance
- Highest SMP scalability (to 128p)
- HP-UX for mission-critical technical computing
- Extensive 64-bit ecosystem (and 32/64-bit on HP-UX)
- Scale-up and scale-out for complex workloads
- HP-UX, Linux & Windows



HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:



RECOMMENDED TRAINING VENUE FOR THE
HP Certified Professional

