



# High Performance Computing, Grids and Clouds: Synergies, Trends and Challenges



Frank Baetke  
(SCI)  
Scalable Computing  
Infrastructure Organization

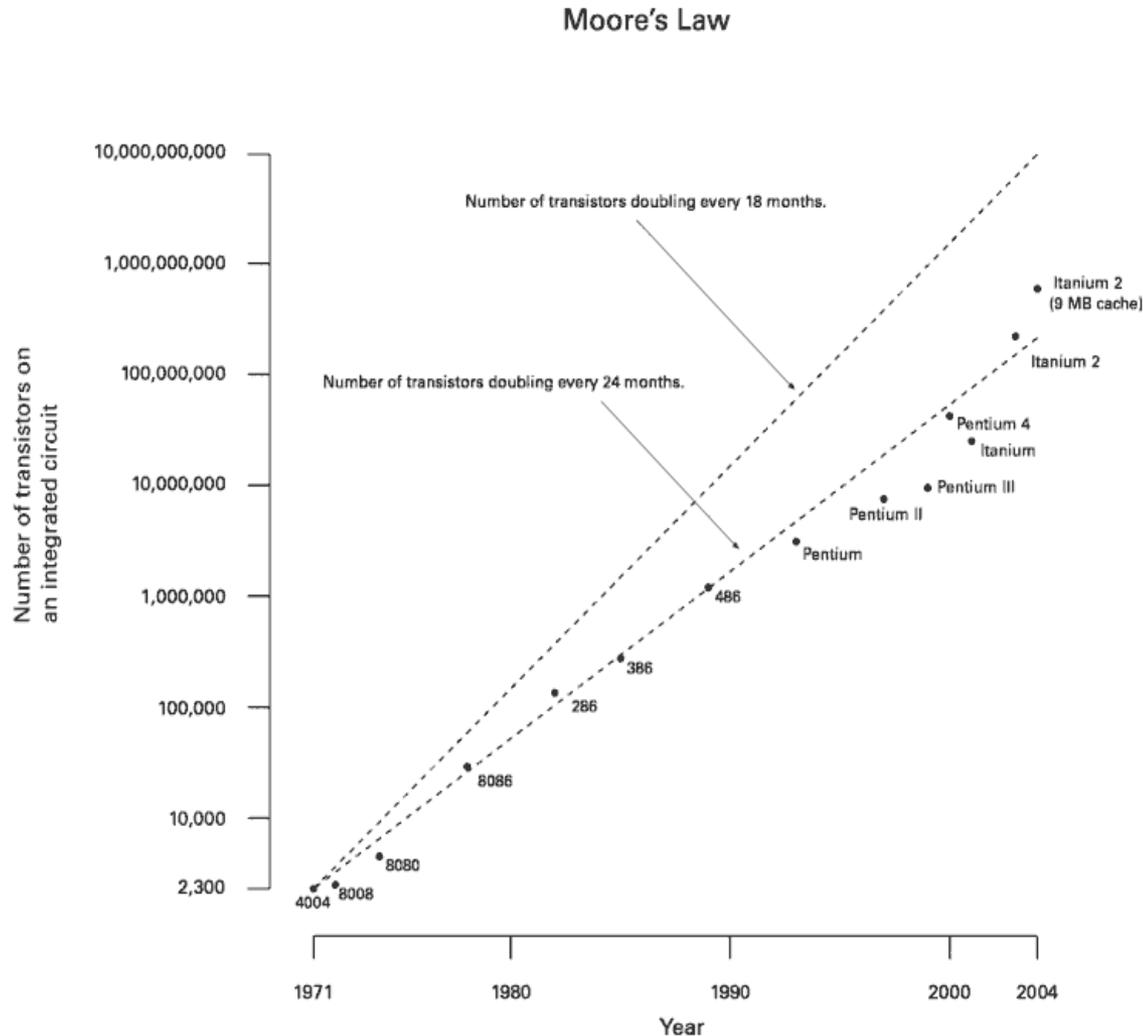
KDM'09 March 13th, 2009

Technology for better business outcomes

# Facts



# Moore's Law → Moore's Cores



# Why the World Goes Multi-core I

- Let **f** be the frequency of the processor and **ipc** the actual or sustained **i**nstructions **p**er **c**lock-tick

then we get: **Perf = f \* ipc**

- To get the frequency **f** up we just need to increase the voltage **V** (and make the chip smaller),

because: **f ≈ V** (in region of interest)

# Why the World Goes Multi-core II

- **Bad news:** the power **E** burned by the chip depends on capacity **C** (transistors etc.), **V** and **f**

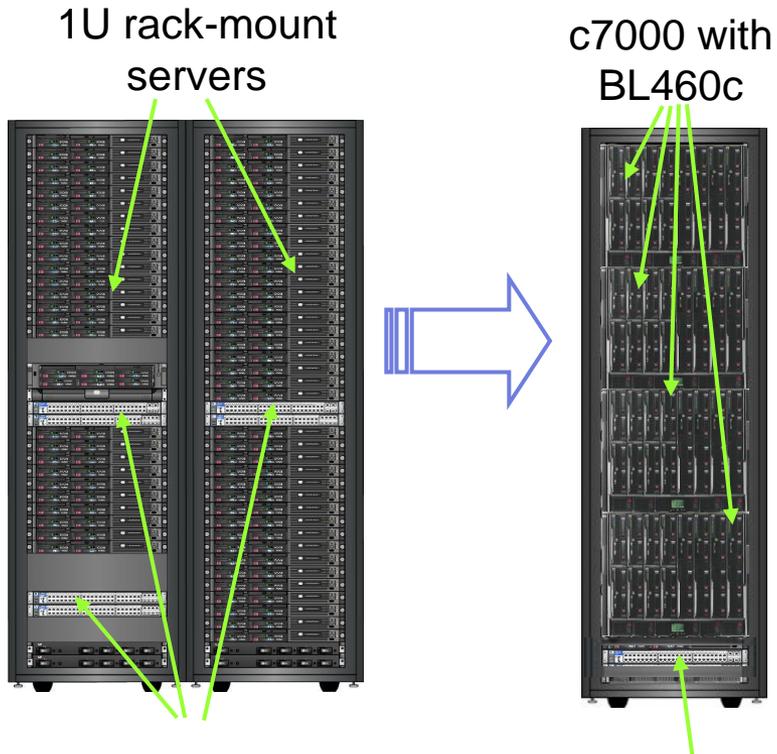
$$E = C * V^2 * f \quad \text{but because } f \approx V$$

$$E = C * V^3$$

- So if we reduce the voltage **V** by 20% we still get 0.8 of the original performance. But **E** goes down dramatically because  $0.8 * 0.8 * 0.8 \approx 0.5!$
- Relative performance  $0.8 * 2 = 1.6$

# HPC Clusters: from rack-mount to blade

Configuration example: cluster of 256 cores



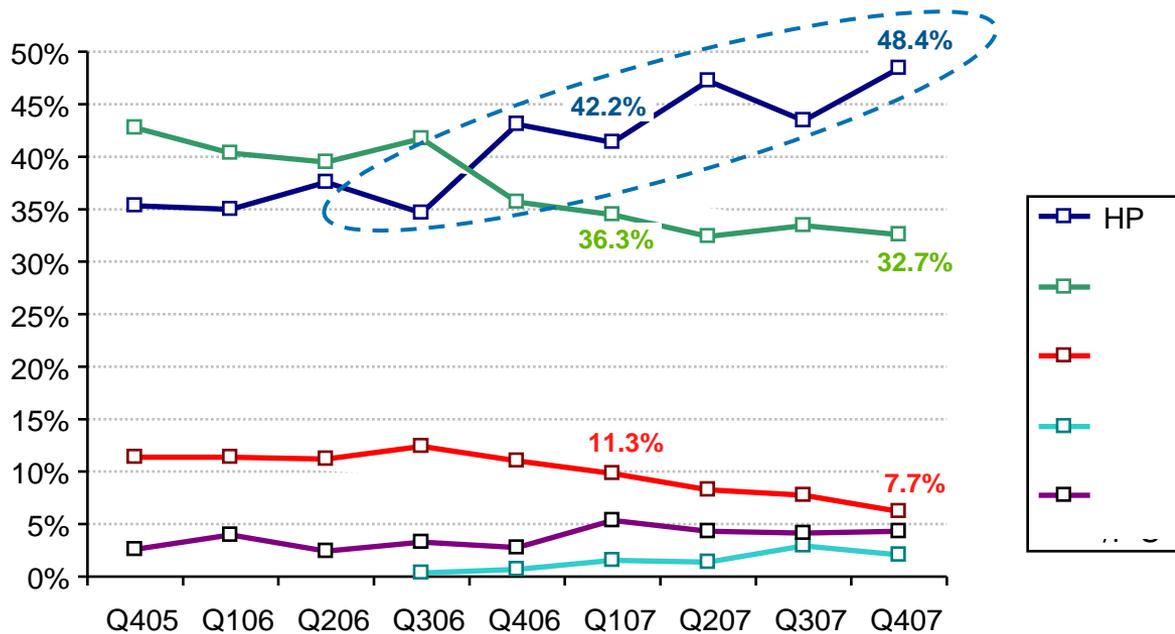
## BladeSystem Advantage

- Floor space: *up to 50% less*
- Power: *up to 40%\* saving*
- Cable complexity: *up to 75% less*
- Power Redundancy: *N+N*

\* Based on internal HP testing of similarly configured rack and blade servers running identical tests

# BladeSystems have been embraced in the marketplace

70% year over year BladeSystem growth!



201 of the top 500 systems in the world run on c-Class  
c-Class in TOP20



HP's Virtual Connect Architecture Wins Product of the Year Award!

**Gartner** Midmarket Summit  
3<sup>rd</sup> "Best in Show" in a row as voted by attending CIO's of midmarket companies



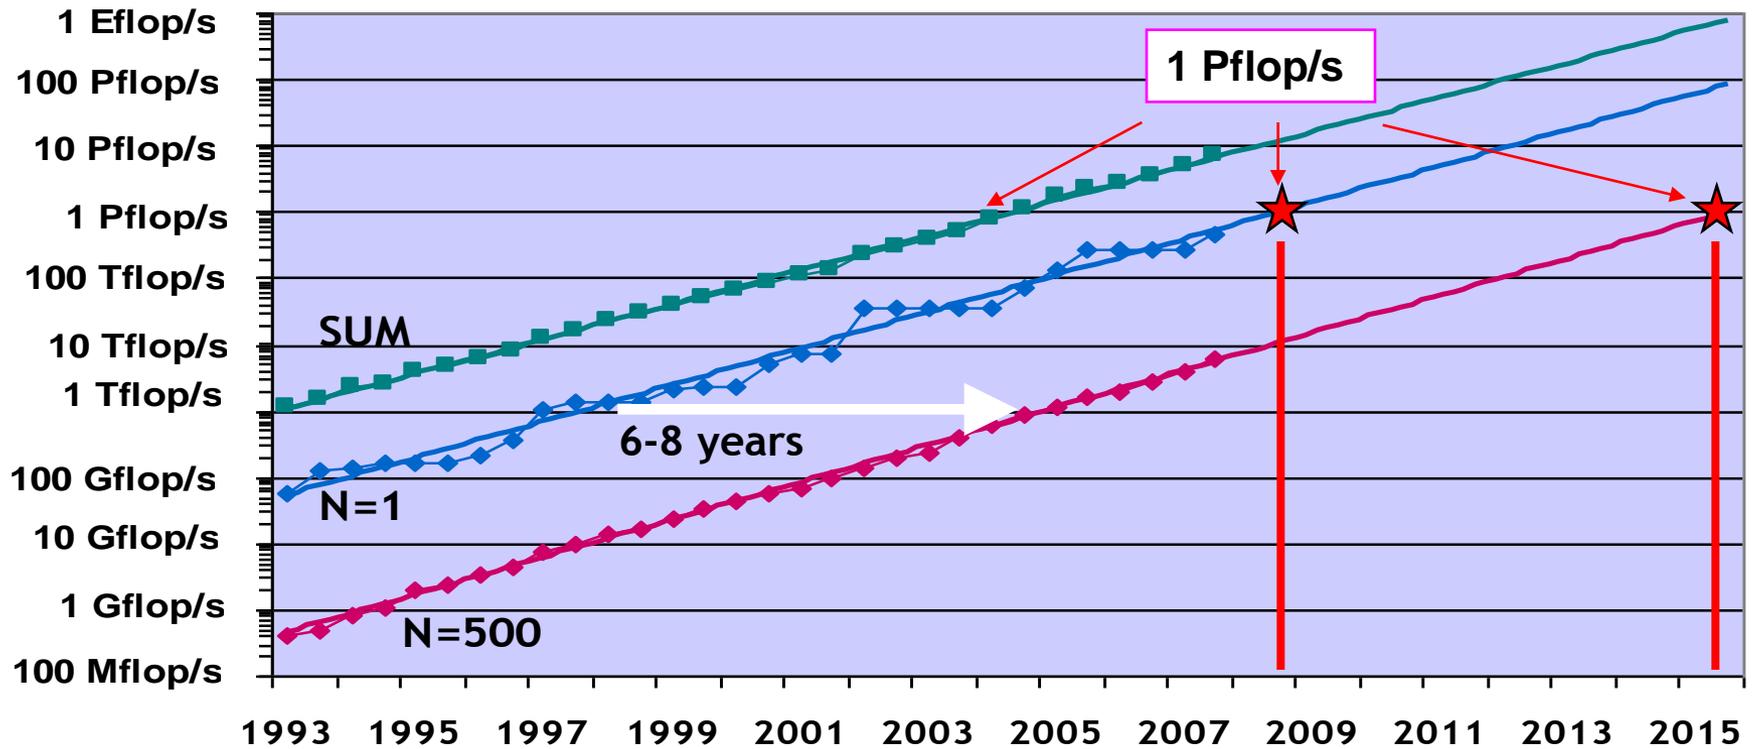
BladeSystem c3000 Most Important Products of 2007



Tech innovator of the Year award Server Hardware category for the second year in a row.



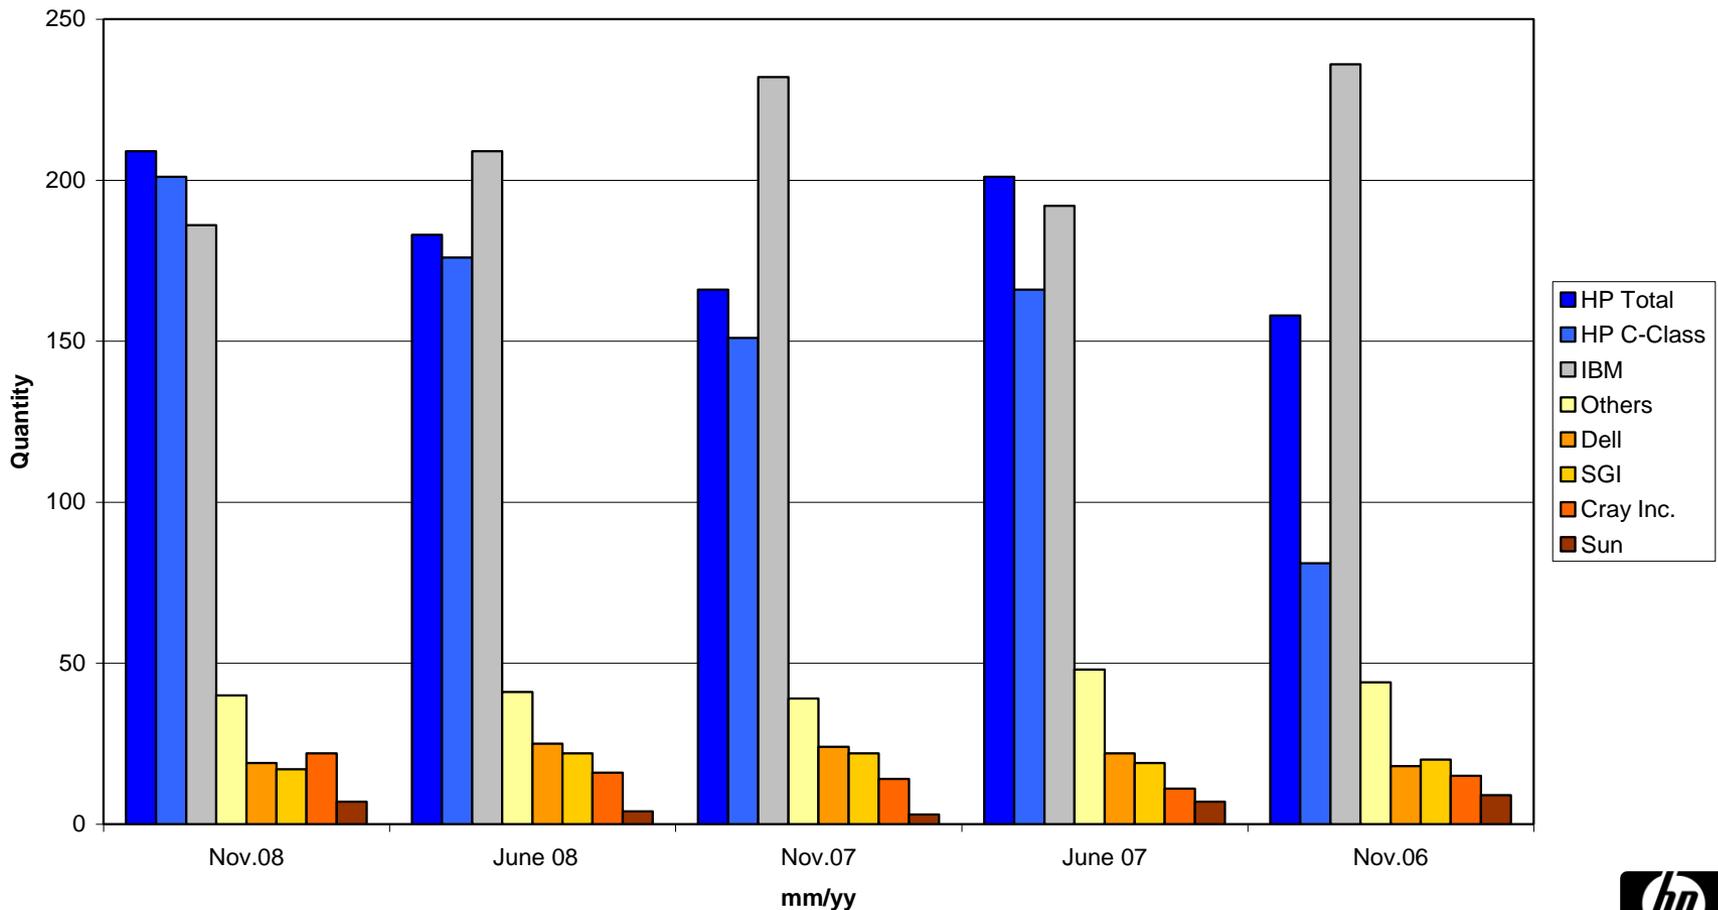
# TOP500 Performance Projections



# TOP 500 list. Nov 08 edition.

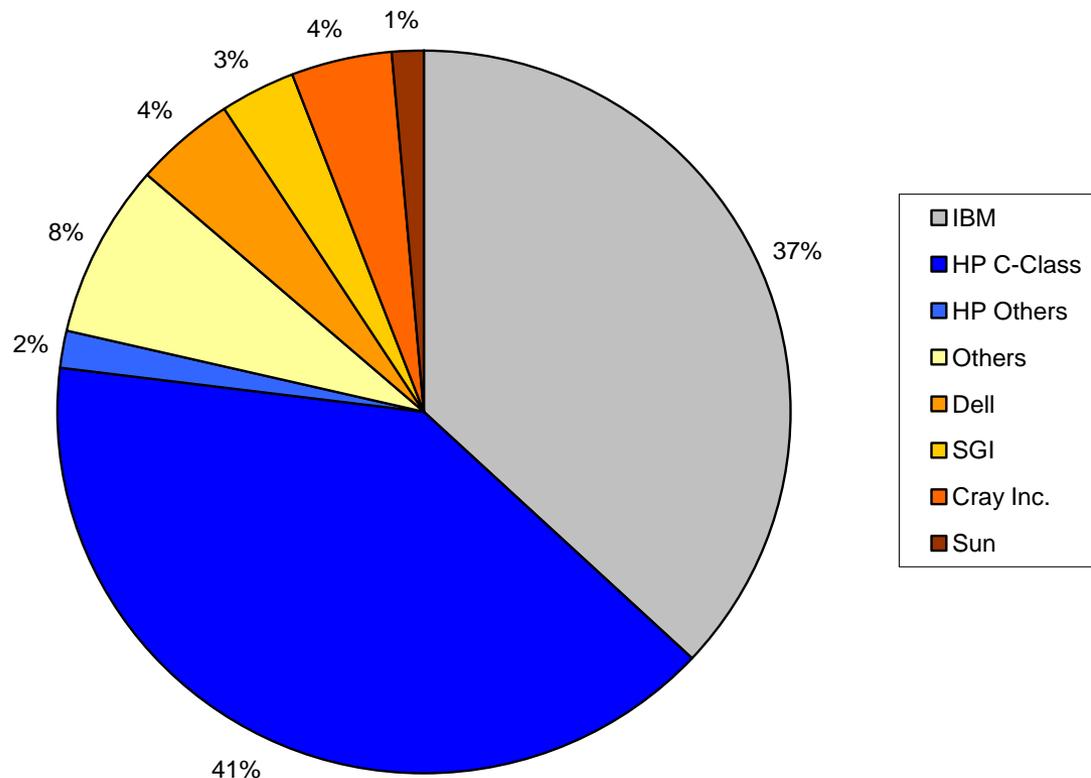
## Trends and entries by system and vendor

TOP500 Vendors Development



# TOP500 list: November 08 edition. Share of *systems* in list by vendors I

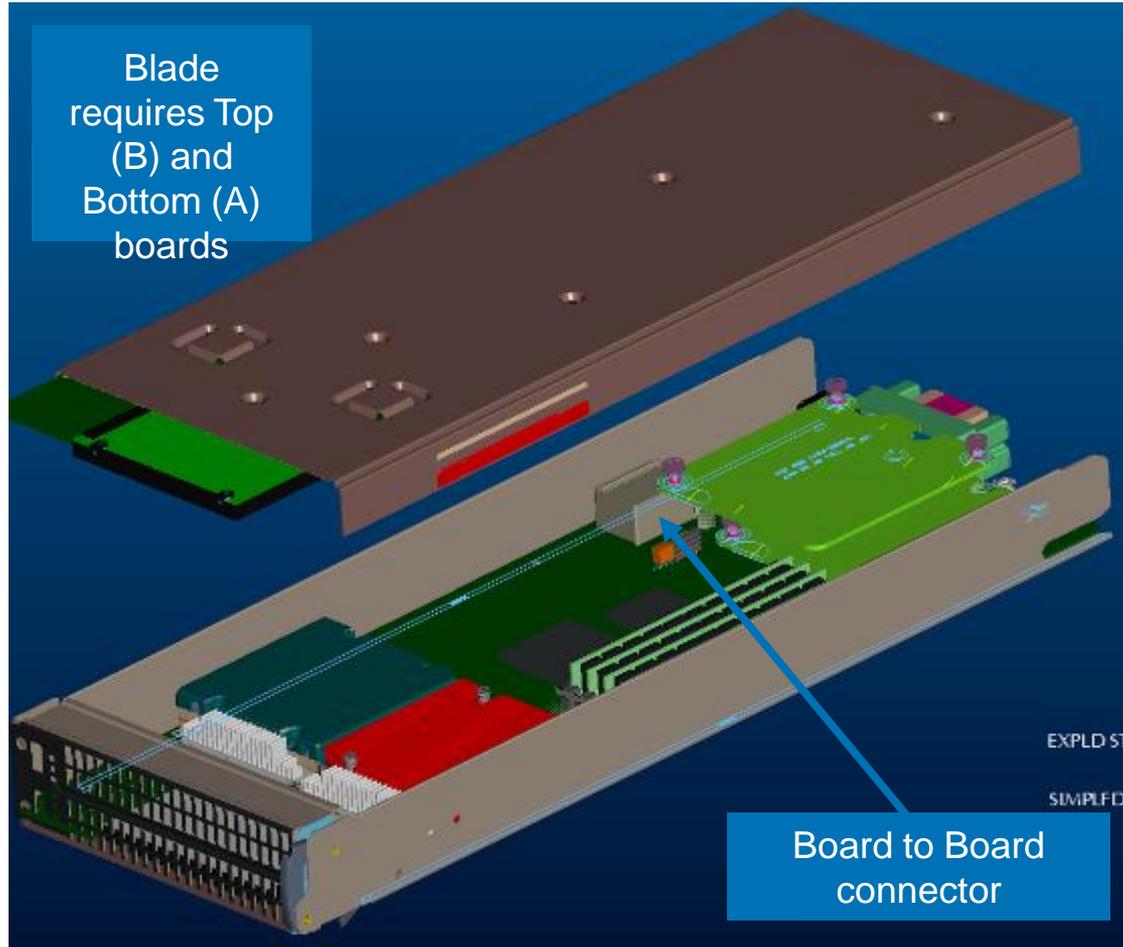
Nov.08



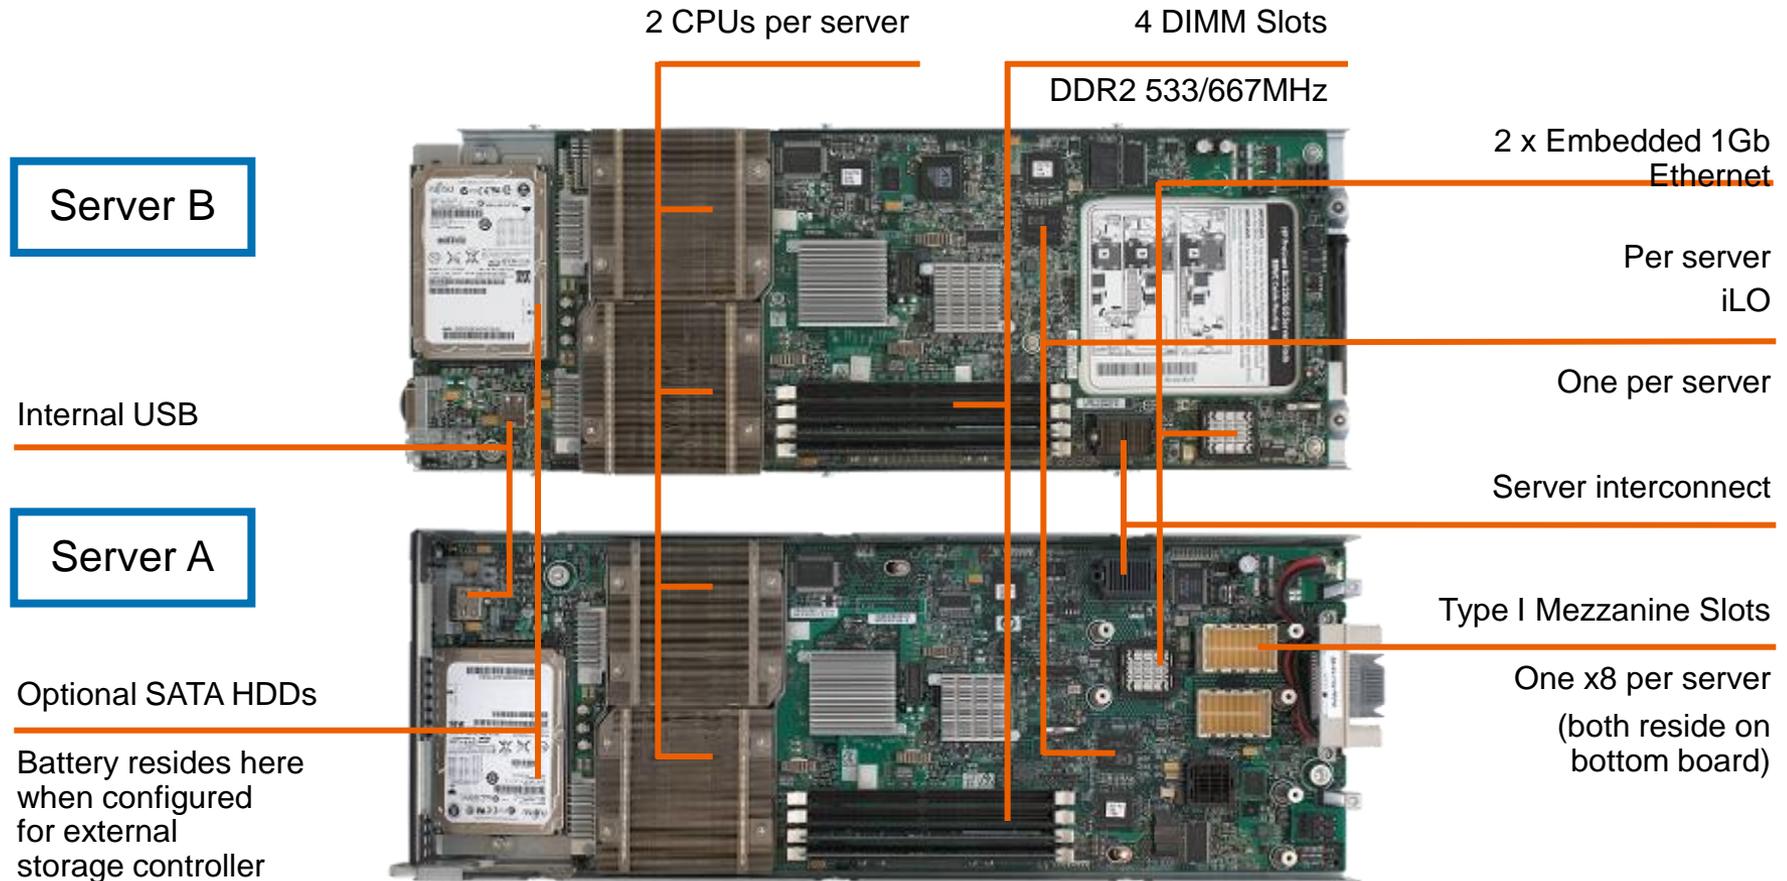
# System Trends



# BL2x220c Assembly



# HP ProLiant BL2x220c G5



# Some interesting data

Eight years ago: ASCI White (IBM)

- #1 on the Top500 in June 2001



Peak performance: 12,288 Gflop/s

Weight: 106 Tons (w/ 160 TB storage)

Power: 3MW

Cost: \$110 million

Today: one rack full of BL2x220c



Peak performance: 12,288 Gflop/s

Weight: ~2000 lbs (~1 ton) – 100x lighter

Power: ~27KW – 100x less power

Cost: ~ \$800K -- more than 100x lower cost

# Cluster Platform Workgroup System (CPWS)

- A winner for the mid-market and volume cluster space
  - 1.5Tflop/s peak performance in 2 sq. ft. floor space
    - 3.0Ghz, 4 flops/cycle, 128 cores
      - $3 \times 4 \times 128 = 1536 \text{Gflop/s} = 1.5 \text{Tflop/s}$
  - No special power & cooling required\*
  - GigE and IB – no network cables



\* Site planning may be required to make sure there is sufficient power draw from the circuitry



# Interconnect Trends: Observations

- Very few proprietary interconnects have a chance to survive – **if any!**
- GigE stays very popular if latency is not an issue - 10 GigE is gaining
- Infiniband has gained significantly
- Future standards based on optical interconnects? Converged Fabrics?

# Software Trends



# Processors and Operating Systems

## Standards and Endangered Species

	LINUX BG	Solaris	AIX	HP-UX	Linux 64	WIN 64
BLUE GENE PowerPC BG	a					
CELL					3x	
SPARC 64		A				
POWER n			A		a	
ITANIUM				A	A	a
x86		a			A	A

# Processors and Operating Systems – Standards and Endangered Species II

	LINUX BG	Solaris	AIX	HP-UX	Linux 64	WIN 64
BLUE GENE PowerPC BG	LE					
CELL					3x ?	
SPARC 64		BE				
POWER n			BE		BE	
ITANIUM				BE	LE	LE
x86		LE			LE	LE

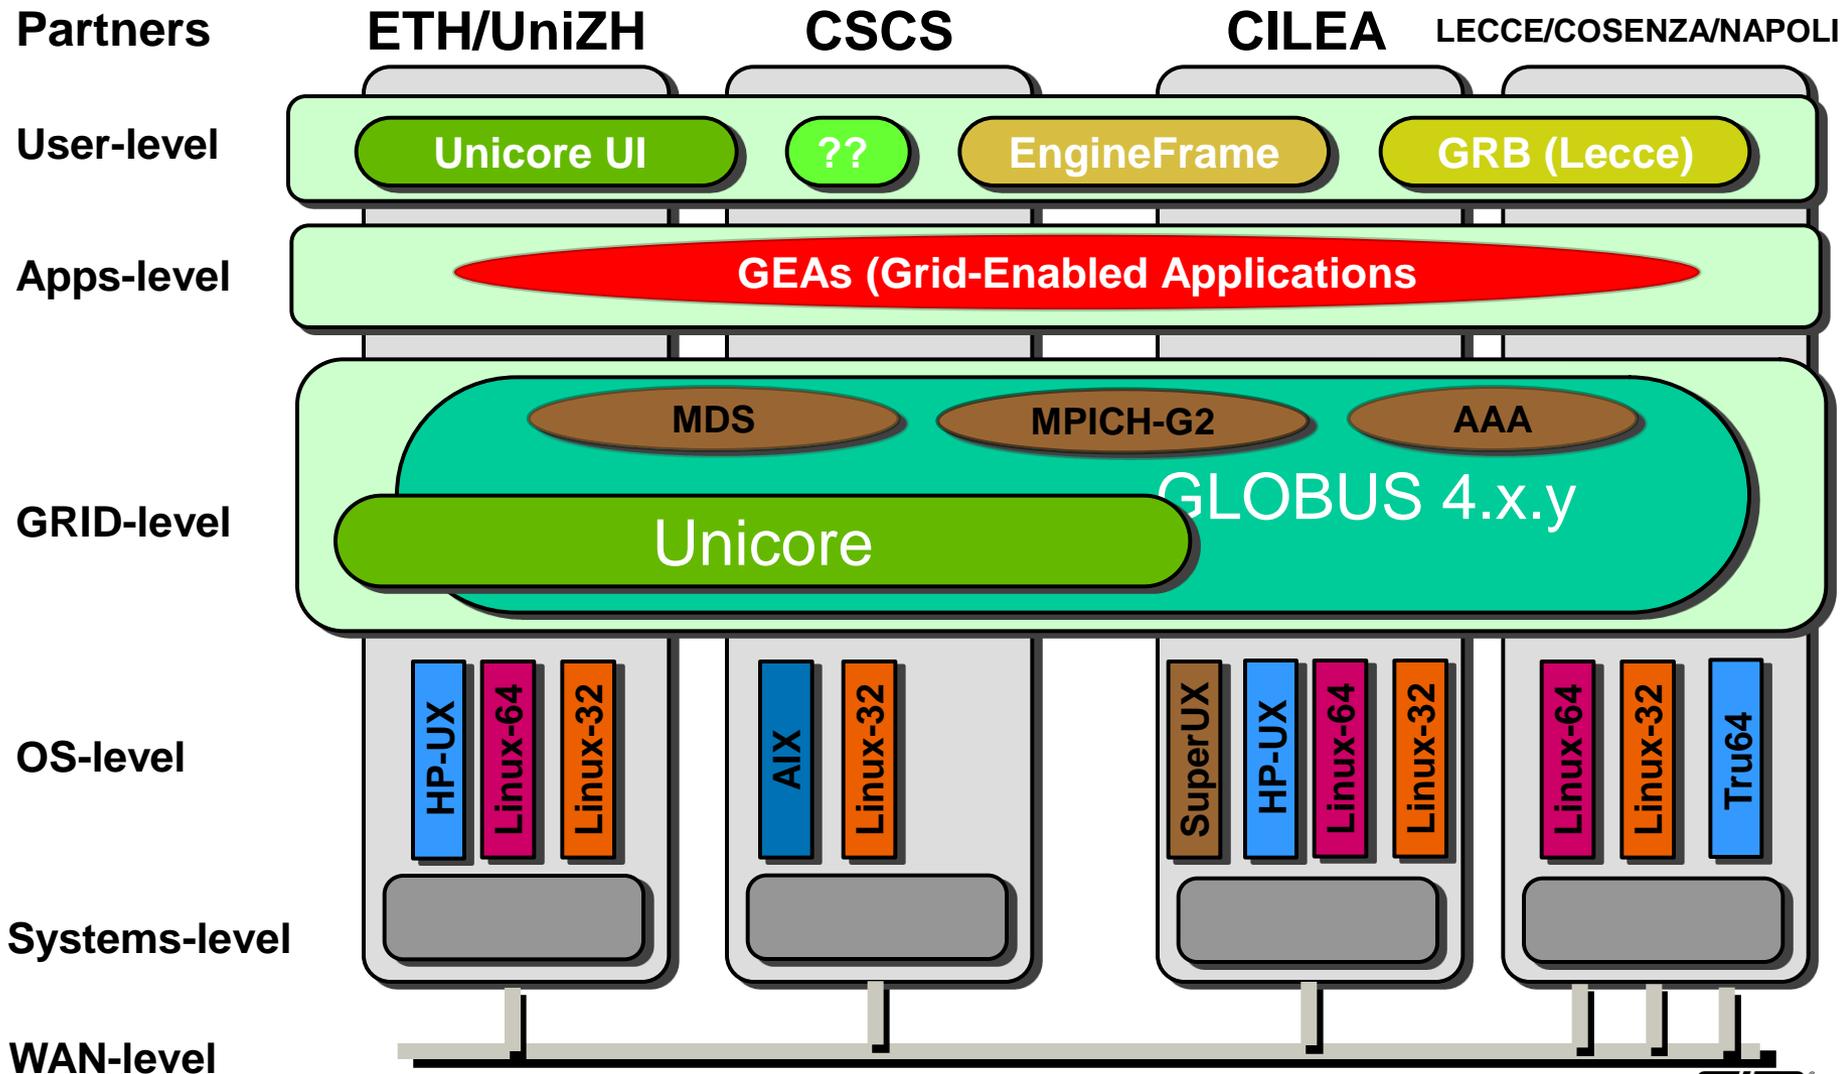
## The Centre of Gravity

	LINUX BG	Solaris	AIX	HP-UX	Linux 64	WIN 64
BLUE GENE PowerPC BG	a					
CELL					3x	
SPARC 64		A				
POWER n			A		a	
ITANIUM				A	A	a
x86		a			A	A

<http://www.sepac-grid.org>



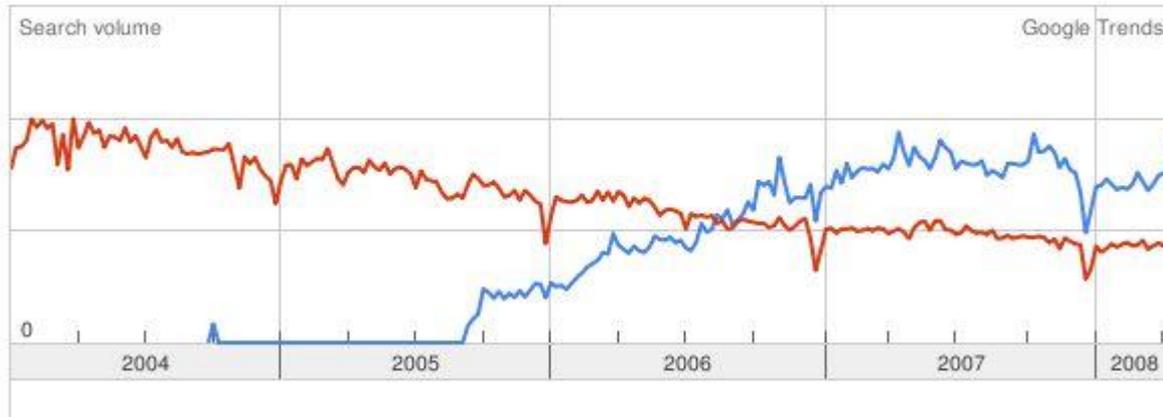
# SEPAC Partner Sites & Logical Grid Structure



# From Google Trends

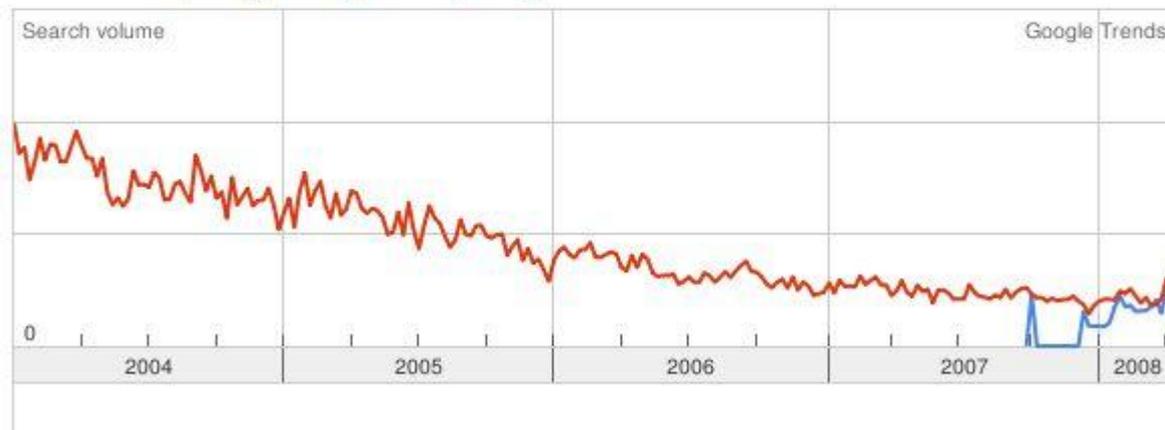
## Trend history

● "web 2.0" ● "web services"



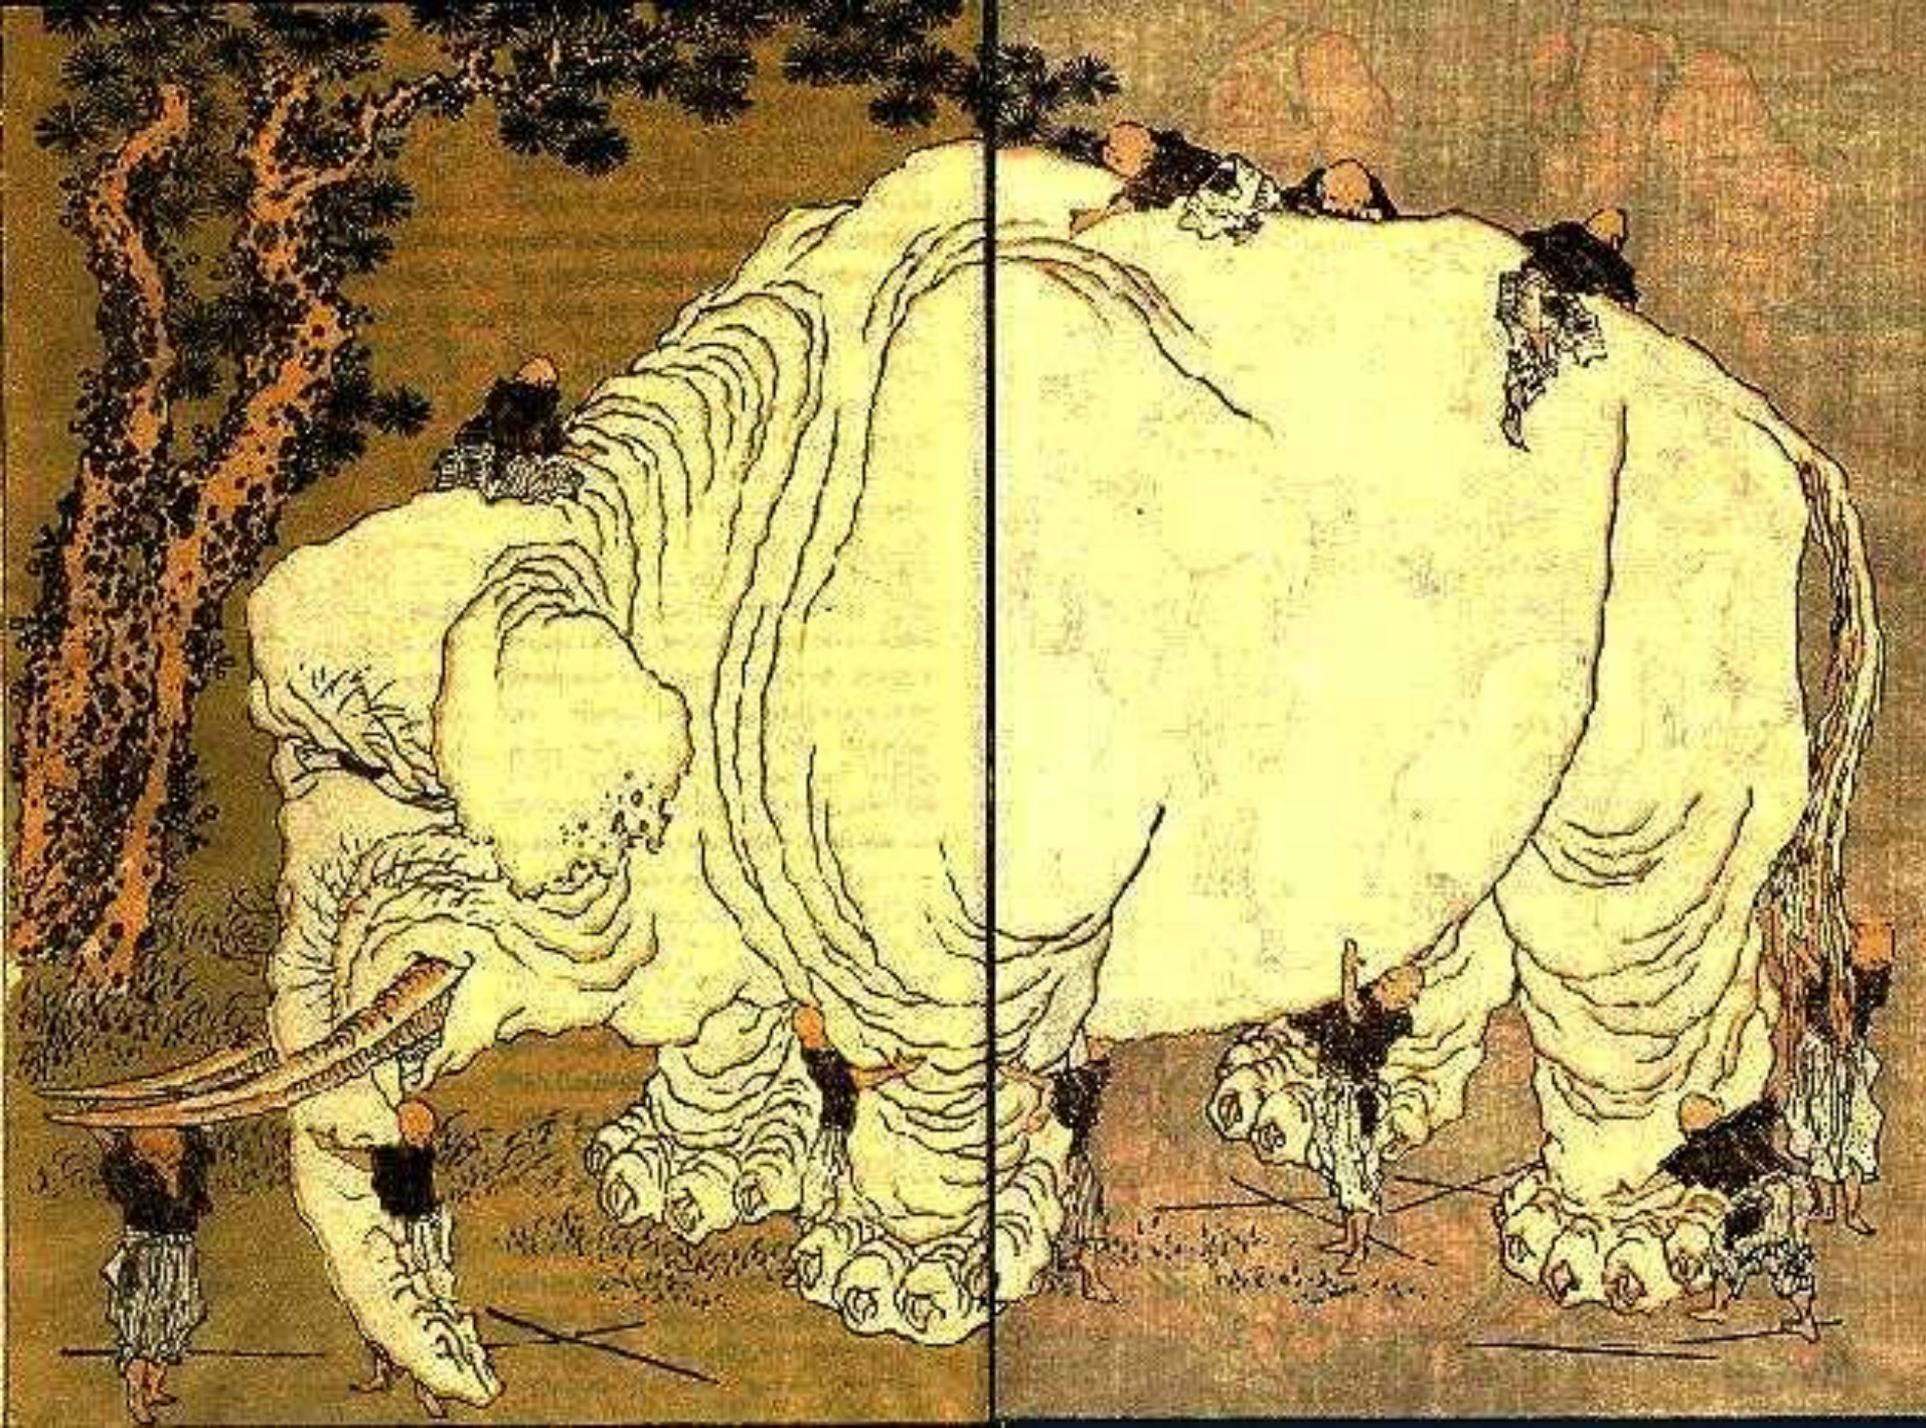
## Trend history

● "cloud computing" ● "grid computing"



# Defining the cloud

- Everybody is trying to figure out what it is (remember grid?)
- Gartner: “Cloud computing is a style of computing where massively scalable IT-related capabilities are provided ‘as a service’ across the Internet to multiple external customers”
- Forrester: “A pool of abstracted, highly scalable, and managed infrastructure capable of hosting end-customer applications and billed by consumption”



# Cloud: the Hype

- “as big as the e-business revolution”
- “will change the nature of computing and scientific research”
- “a disruptive technology”
- The Gartner hype-cycle:
  - Technology Trigger,
  - Peak of Inflated Expectations,
  - Trough of Disillusionment,
  - Slope of Enlightenment,
  - Plateau of Productivity

# Grid and Cloud Today

## GRID

- open standards (OGF ...)
- publicly funded & operated (slow evolution)
- no central management
- interoperability important
- geographically distributed; locally owned and managed
- share (usually modest) local resources
- scientific research, high-end users

# Grid and Cloud Today

## GRID

- open standards (OGF ...)
- publicly funded & operated (slow evolution)
- no central management
- interoperability important
- geographically distributed; locally owned and managed
- share (usually modest) local resources
- scientific research, high-end users

## CLOUD

- no standardized interfaces
- privately funded & operated (fast evolution)
- managed by a single entity
- no interoperability
- geographically distributed; centrally owned and managed
- make huge systems available
- enterprise applications, information processing, data mining

# Cloud Computing Vendors

- <http://www.johnmwillis.com/mysql/cloud-vendors-a-to-z/>
- 3Tera, Adobe Air, Akamai, Amazon EC2, Amazon S3, Amazon SimpleDB, Apache CouchDB, Apache Hadoop, Areti Internet, Box-Net, Cassatt Corporation, Citrix (XenSource), CohesiveFT, Dell DCS, Elastra, EMC Mozy, Enki, Enomoly, Enomoly ElasticDrive, EnterpriseDB, Flexiscale, Fortress ITX, Google Apps, HP AlaaS, IBM Blue Cloud, iCloud, Joyent, JungleDisk, Layered Technology, LongJump, Microsoft SSDS, MorphExchange, Mosso, Rackspace, Rightscale, Salesforce.com, Sun Caroline, Sun MySQL, Terremark, VMWare

# Challenges



“Mainstream computing needs parallelism for performance. This is new.”

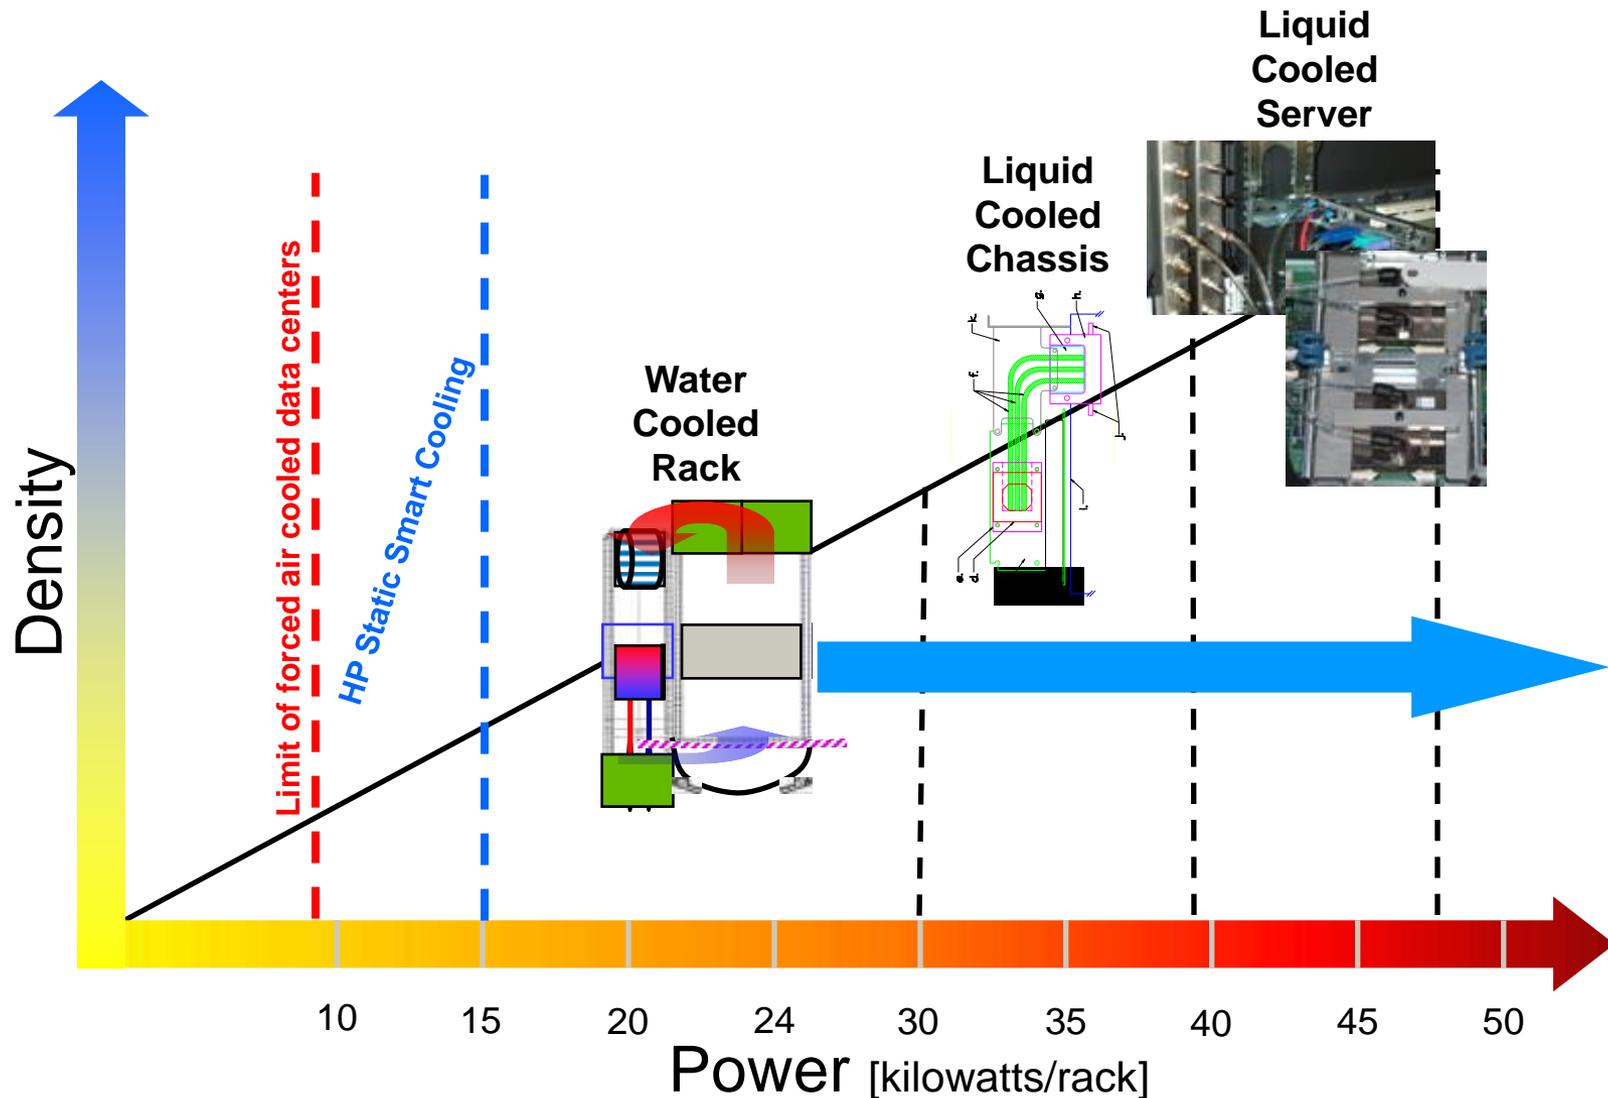
“A lot of very ordinary programmers will have to write parallel programs. Very few experts currently know how to do that.”

Hans Boehm, Advanced Architecture Lab,  
HP Labs

# The Other Challenge ...

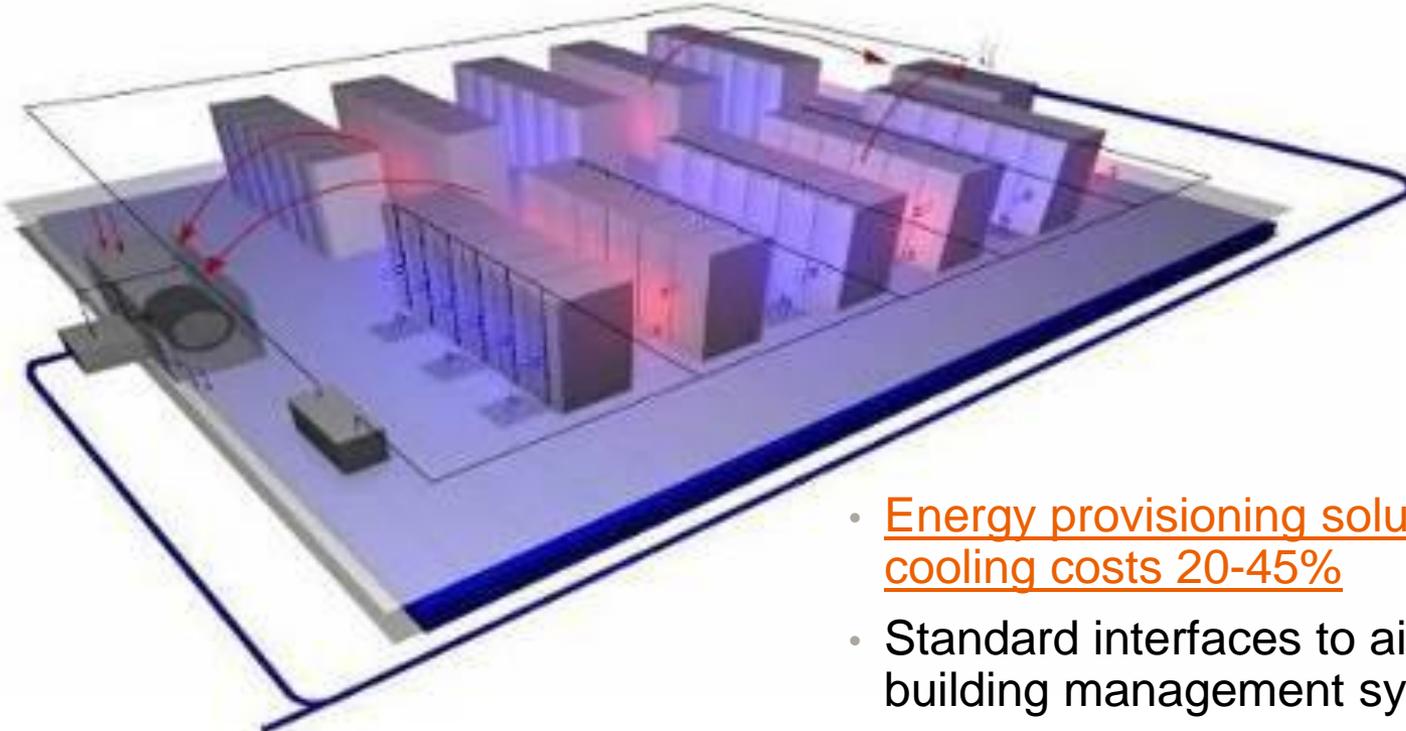


# Rack Density vs. Alternate Cooling



# HP Dynamic Smart Cooling

*Bridging Facilities and IT to realize Adaptive Infrastructure*



“Dynamic Smart Cooling is the most remarkable development for data center critical support systems.”

-Peter Gross, CEO and CTO, EYP Mission

- Energy provisioning solution can reduce cooling costs 20-45%
- Standard interfaces to air-conditioning and building management systems
- Pervasive thermal sensing grid at rack level with zoning capabilities at room level
- Easy to retrofit or spec for new construction applications
- Available for commissioning *NOW!*
- Compatible with 3<sup>rd</sup> party gear

# SCI Infrastructure HP POD (Container)

## Industry-standard Flexibility

22 x 50U, 19" full-depth industry-standard racks support HP, Dell, IBM, Sun, Cisco, etc.

## Best-in-class Density

Support for 3,520 compute nodes, 12,000 LFF drives, or any combination

## Shipped in 6 weeks, deployed WW

Pre-integrated, configured and tested before shipment; shipped in six weeks from order.

## Energy Effectiveness

PUE ratio <1.25 (1.07 excluding chiller)

## Infrastructure Services Portfolio

Full lifecycle support services combining technology and facilities expertise



# Summary



# Convergence of Massive Scale-Out

## HPC

- Engineering & Geo-Sciences
- Life & Materials Sciences
- Defense/Security
- Scientific Research

## Enterprise Scale-Out

- Financial Analytics
- Real-time Trading
- Schedule Optimization

## Web 2.0

- Internet Commerce
- Interactive Media
- Online Gaming



## Emerging business models

Best performance = faster time to market

## New metrics

Best performance per watt

Best performance per sq. ft.

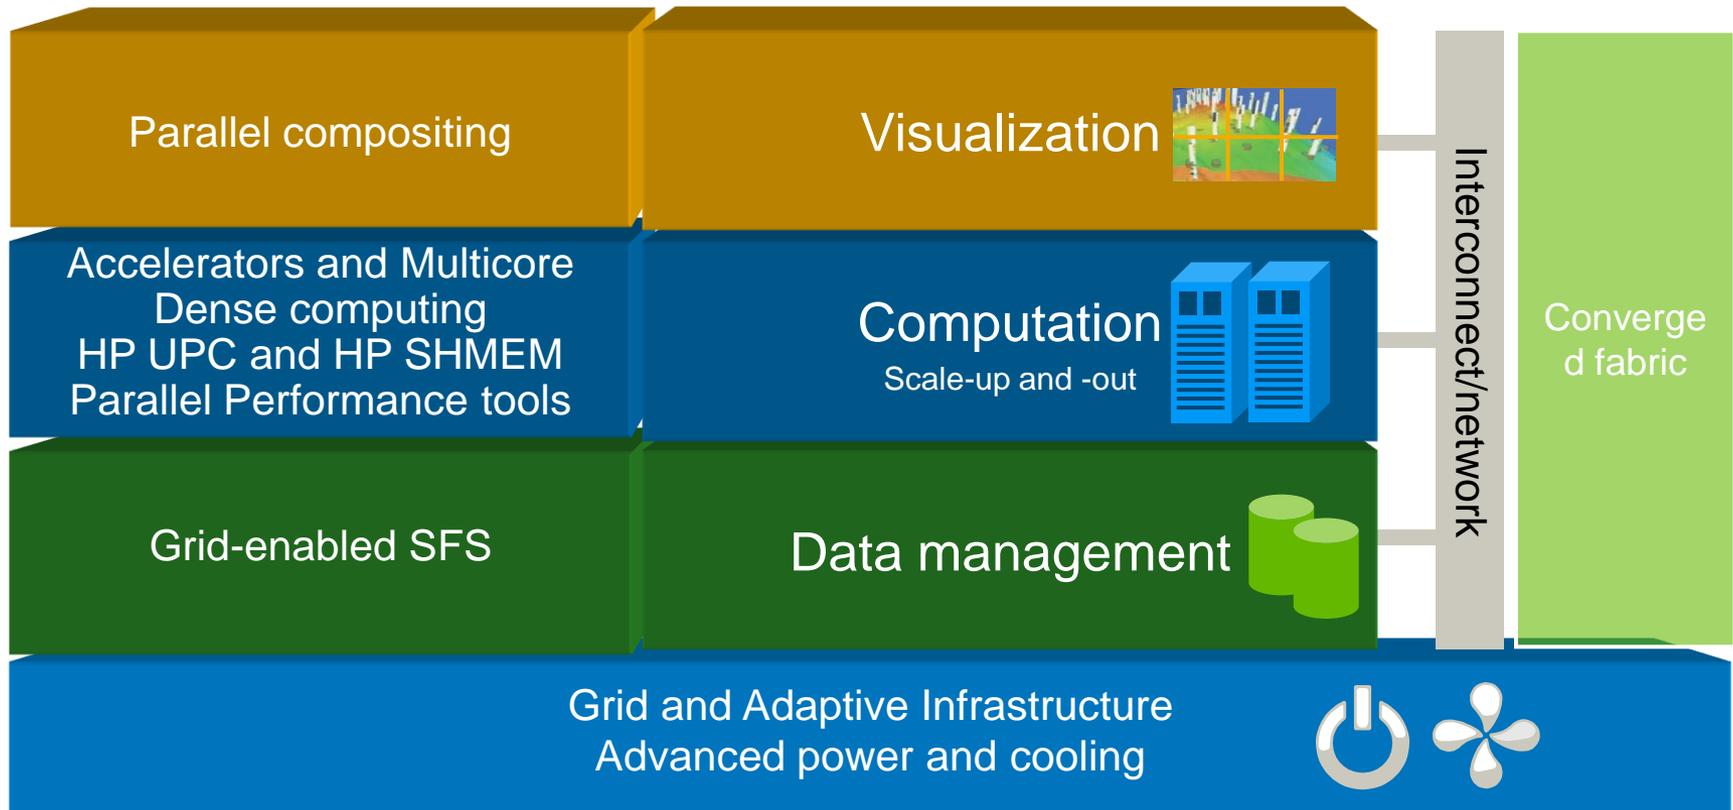
## Extreme pain points

Data center constraints of power, cooling, space, manageability & automation of dynamic workloads

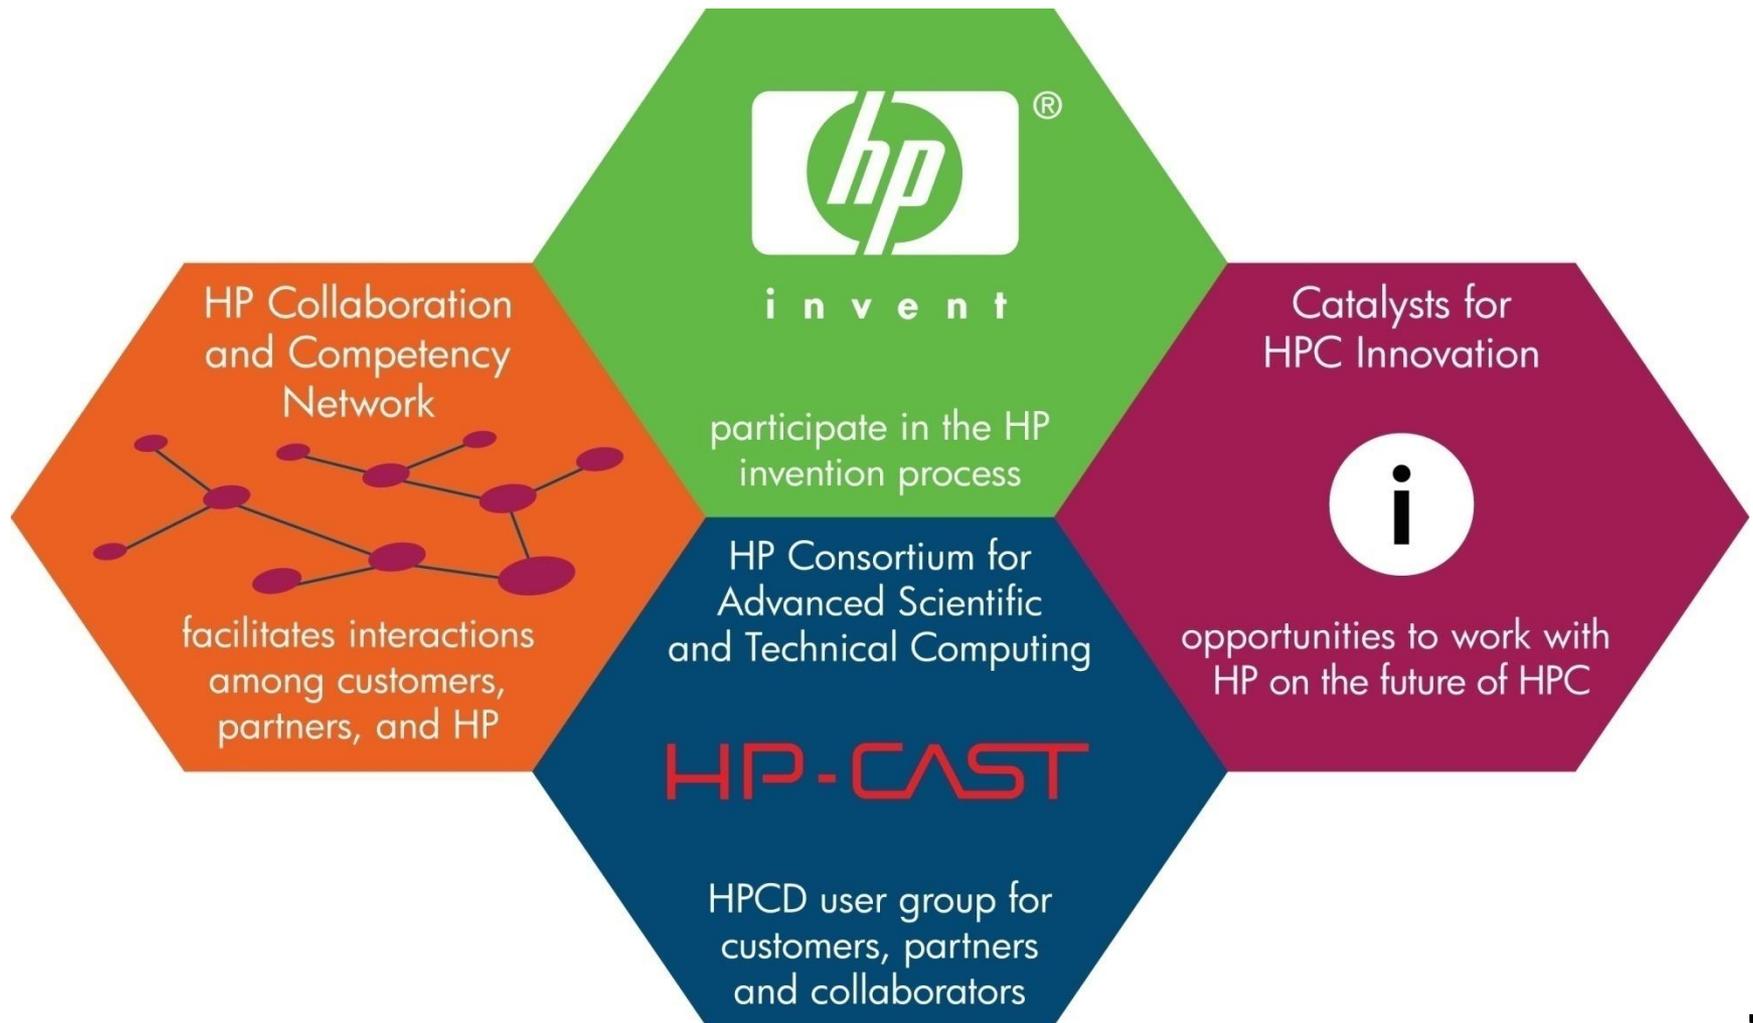
## Customer Engagement

Multiple procurement models, requiring optimized supply chains & unique products and services

# Extending the vision: powerful catalysts for HPC *innovation*



# Working collaboratively to accelerate the pace of innovation



# HP-CAST 12 Status Update

- SCI's 12th edition of the world-wide user group conference will take place in Madrid, May 10-13.
- Integrates previous EMEA „Executive Forum“.
- Hosts „Nordic and Hispanic Regional Users Group.
- First calls for papers issued, first customer talks accepted – for more details, see: [www.hpcast.org](http://www.hpcast.org)
- Registration will open in March. Venue: Sol Melia Princesa, low registration fee (as always) €229.
- **Please contact me directly if you are planning to attend and/or present a paper or to participate in a tutorial**

# HP-CAST 12 = 5 1/2 Days of Technology, Customer and Partner Updates

- Sunday, May 10th Eight deep dive technical tutorials (4 tracks): Interconnects, accelerators, administration, cloud computing, PODs etc.
- Monday, May 11th Invited customer keynotes, SCI product and roadmap updates
- Tuesday, May 12th „Interconnect Focus Session“. Customer break-out sessions
- Wednesday, May 13th „Datacenter Focus Session“. Customer issues, collaboration and catalyst meeting, Large System SIG etc.

## Announcements.

HP Consortium for Advanced Scientific and Technical Computing  
User Group Meeting  
Scalable Computing Infrastructure (SCI) Organization  
Venue: Center Madrid  
May 10<sup>th</sup> – 13<sup>th</sup>, 2009 – Madrid, Spain  
**HP-CAST 12 / Executive Forum & NTIG (Nordic Technical Interest Group) & HP-CAST IBERICA**  
Draft Agenda Version 1.3  
Sunday, May 10<sup>th</sup> – Registration and Tutorials

10:00 – 18:00	Registration	
12:30 – 15:00	<b>Tutorial_A1</b> Cloud Computing Basics	
12:30 – 15:00	<b>Tutorial_B1</b> Multi-core Processor Technologies	
12:30 – 15:00	<b>Tutorial_C1</b> Performance / Interconnect Topologies	
12:30 – 15:00	<b>Tutorial_D1</b> t. b. d.	
15:00 – 15:30	Break	
15:30 – 18:00	<b>Tutorial_A2</b> Cloud Computing Implementations	
15:30 – 18:00	<b>Tutorial_B2</b> Accelerator Technologies	
15:30 – 18:00	<b>Tutorial_C2</b> Datacenter Containers (PODs)	
15:30 – 18:00	<b>Tutorial_D2</b> t. b. d.	
19:00	HP-CAST Reception	All Attendees

HP-CAST 12 & Executive Forum

Monday May, 11<sup>th</sup> - Conference Day 1

8:00 – 17:00	Registration	
08:00 – 08:15	HP-CAST Board Opening Remarks HP-Liaison President	Frank Beetke, HP SCI Rudolf Lehner, KIT SCC
08:15 – 09:00	Welcome Keynote: Strategies and Directions of HP's Scalable Computing Organization	Christine Martino Curt Belusar Scott Misage
09:00 – 09:30	Invited Academic Keynote: HPC & Scalable Computing Trends	N. N.
09:30 – 10:00	Invited Governmental Keynote: Towards European PFLDP-Sites	N. N.
10:00 – 10:30	Invited Industrial Keynote: Industrial Cloud Computing Trends	N. N.
10:30 – 11:00	Break	
11:00 – 12:30	New and Strategic Site Updates	SCI Customers
12:30 – 13:30	Lunch	
13:30 – 14:00	AMD Roadmap Update (NDA)	N. N.
14:00 – 14:30	Intel Roadmap Update (NDA)	N. N.
14:30 – 15:30	HP-SCI Platform & Technology Trends (Part 1) Products & Roadmaps in Detail	Ed Turket et al.
15:30 – 16:00	Break	
16:00 – 17:00	HP-SCI Platform & Technology Trends (Part 2) Products & Roadmaps in Detail	Steve Cummings et al.
17:00 – 17:30	Toward Baseline Computing: Challenges and Outlook	N. N. HP - Labs
17:30 – 18:00	Transformations and Trends in HPC: A Critical Review	Richard Kaufman, HP SCI
20:00	Gala Dinner	

HP-CAST 12 / Executive Forum  
Tuesday May, 12<sup>th</sup> - Conference Day 2

Focus Session 1: Advances in Interconnects and Interconnect Topologies

8:00 – 8:30	Invited Lecture: High Performance Wide-Area Networks – the Foundation for Clouds	N. N.
8:30 – 10:00	Implementations of High Performance Interconnects	Short Updates from Interconnect Vendors
10:00 – 10:30	Panel Discussion	Chair: Moray McLaren All Speakers
10:30 – 11:00	Break	
11:00 – 12:30	High Performance and Scalable Computing in the North: Concepts, Sites, Architectures, Trends	NTIG - Nordic Technical Interest Group Speakers
12:30 – 13:30	Lunch	
12:30 – 15:00	<b>Track 1: Lectures</b> Focus on Cloud & Grid Computing	
12:30 – 15:00	<b>Track 2: Lectures</b> Academic Sites – Practice & Experience	
12:30 – 15:00	<b>Track 3: Lectures</b> NTIG – Practice and Experience	NTIG Speakers
15:00 – 15:30	Break	
15:30 – 18:00	<b>Track 4: Lectures</b> Focus on Web2 Hosting	
15:30 – 18:00	<b>Track 5: Lectures</b> Industrial Sites – Practice & Experience	
15:30 – 18:00	<b>Track 6: HP-CAST IBERICA</b> Regional Meeting (in Spanish)	
19:00	HP-CAST Dinner Event	All Attendees

HP-CAST 12 / Executive Forum  
Wednesday May, 13<sup>th</sup> - Conference Day 3

Focus Session 2: Datacenter Design: New Trends and Challenges

8:30 – 10:30	Multiple Invited Lectures and Partner Presentations Cooling and Design Challenges of Large Data Centers Implementations Revised in Detail	N. N.
10:30 – 11:00	Break	
11:00 – 11:30	Customer Collaboration and Support Session	Scott Misage et al. HP SCI
11:30 – 12:00	Future Structure of HP-CAST HP-CAST Elections	HP-CAST Board et al.
12:00 – 12:15	HP Wrap-up and Outlook	Christine Martino, HP SCI
12:15 – 12:30	Logistical Announcements End of Conference Segment Announcement of HP-CAST 13	Rudolf Lehner, KIT SCC Frank Beetke, HP SCI
12:30 – 13:30	Lunch	
13:30 – 15:00	Large System SIG	Gary Skousson, PNNL Scott Misage, HP SCI
15:00 – 15:30	Break	
15:30 – 17:00	System Management Collaboration Session	Las Kramer, SCI and CCN Advocates & Members
15:30 – 17:00	CCN Collaboration Session	N. N.
15:30 – 17:00	Catalyst Session	N. N.
15:30 – 17:00	t. b. d.	N. N.

End of HP-CAST 12

Thank You

