FUJITSU

shaping tomorrow with you

Fujitsu and the HPC Pyramid

Wolfgang Gentzsch
Executive HPC Strategist (external)
Fujitsu Global HPC Competence Center

INTERNATIONAL
SUPERCOMPUTING CONFERENCE

June 20th, 2012

"Fujitsu's objective is to contribute to a prosperous future through the benefits of supercomputers. To accomplish this, we need the collaboration of researchers & scientists worldwide."

Masahiko Yamada
President of Fujitsu's Technical Computing Solutions Unit

HPC matters …

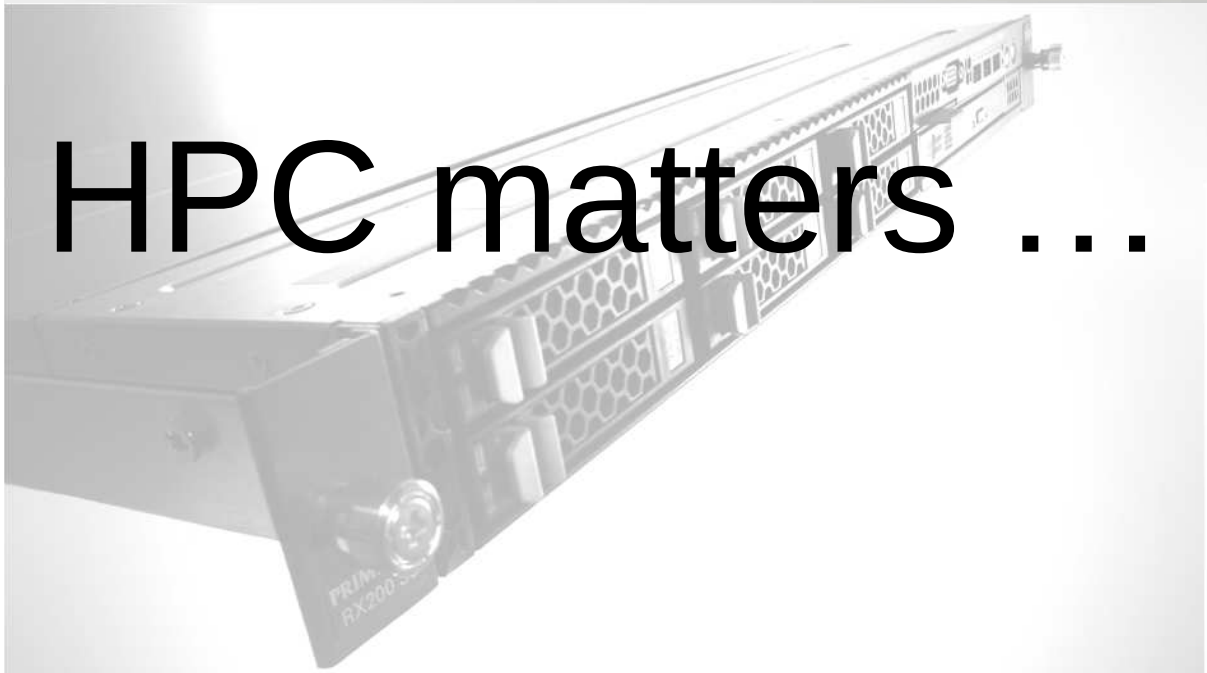

… it matters for Science and Industry

FUJITSU

- 2011 worldwide HPC revenue share by segment

Source: IDC, April 2012

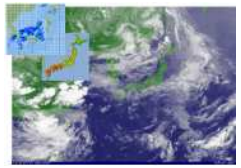
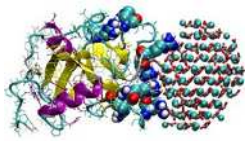

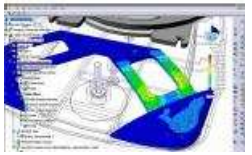
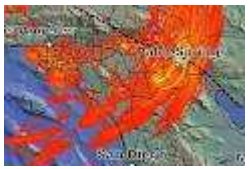

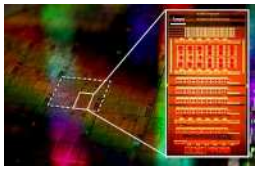
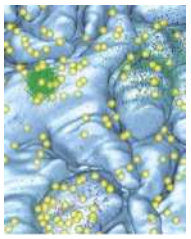

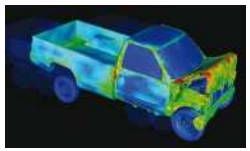


The HPC Pyramid

FUJITSU

Matching Science and Industry needs to the HPC Pyramid:

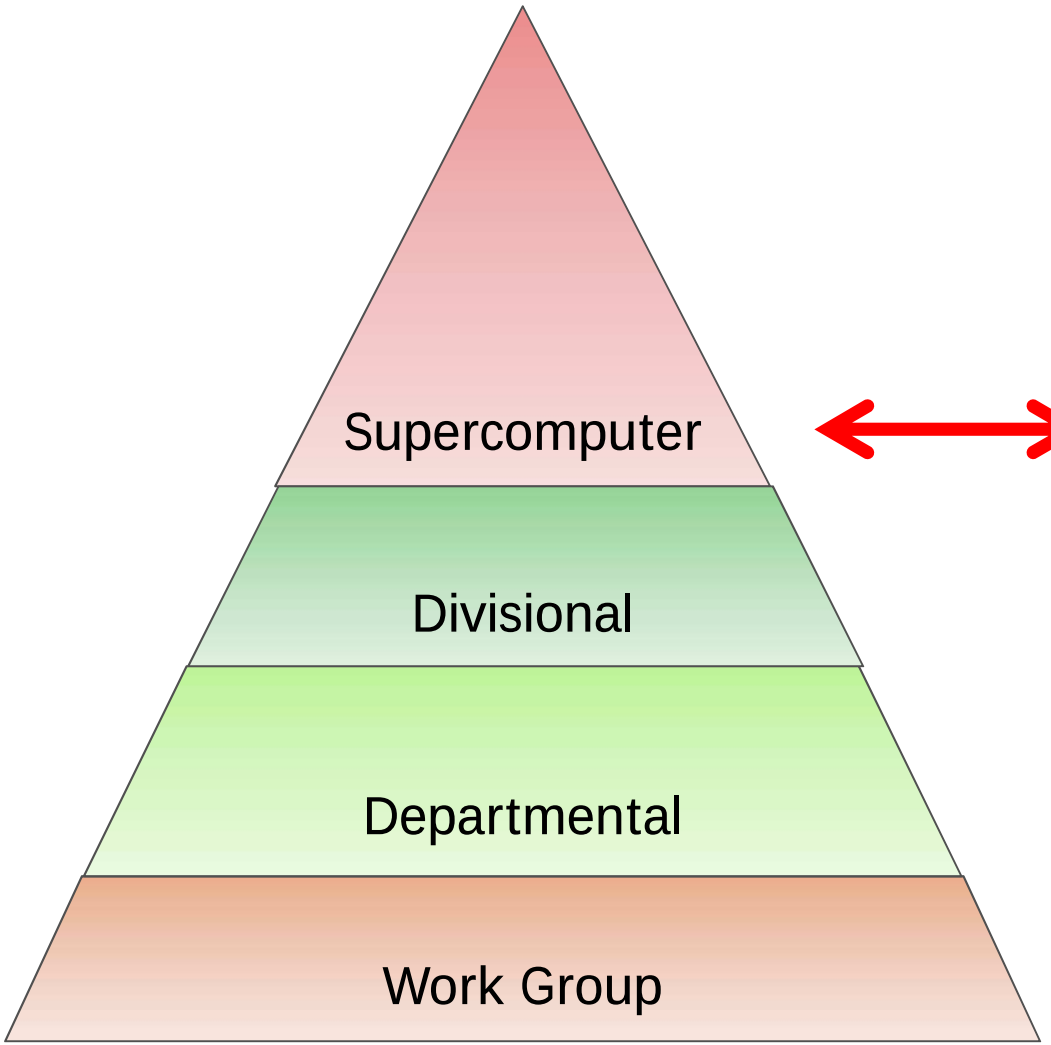

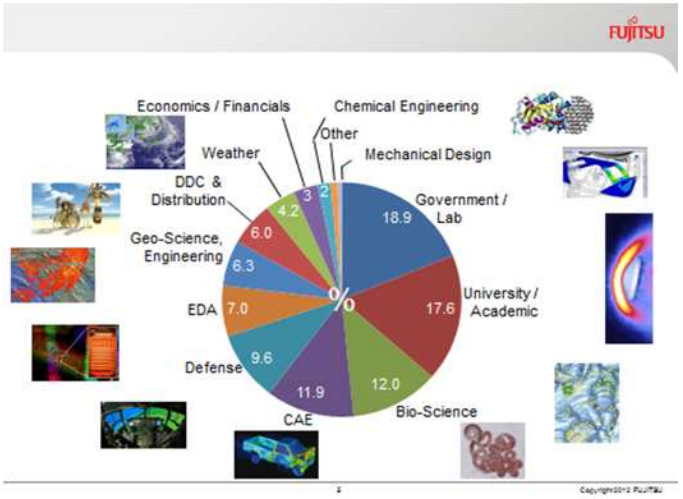

The HPC Pyramid

FUJITSU

Matching Science and Industry needs to the HPC Pyramid:

HPC is becoming mainstream

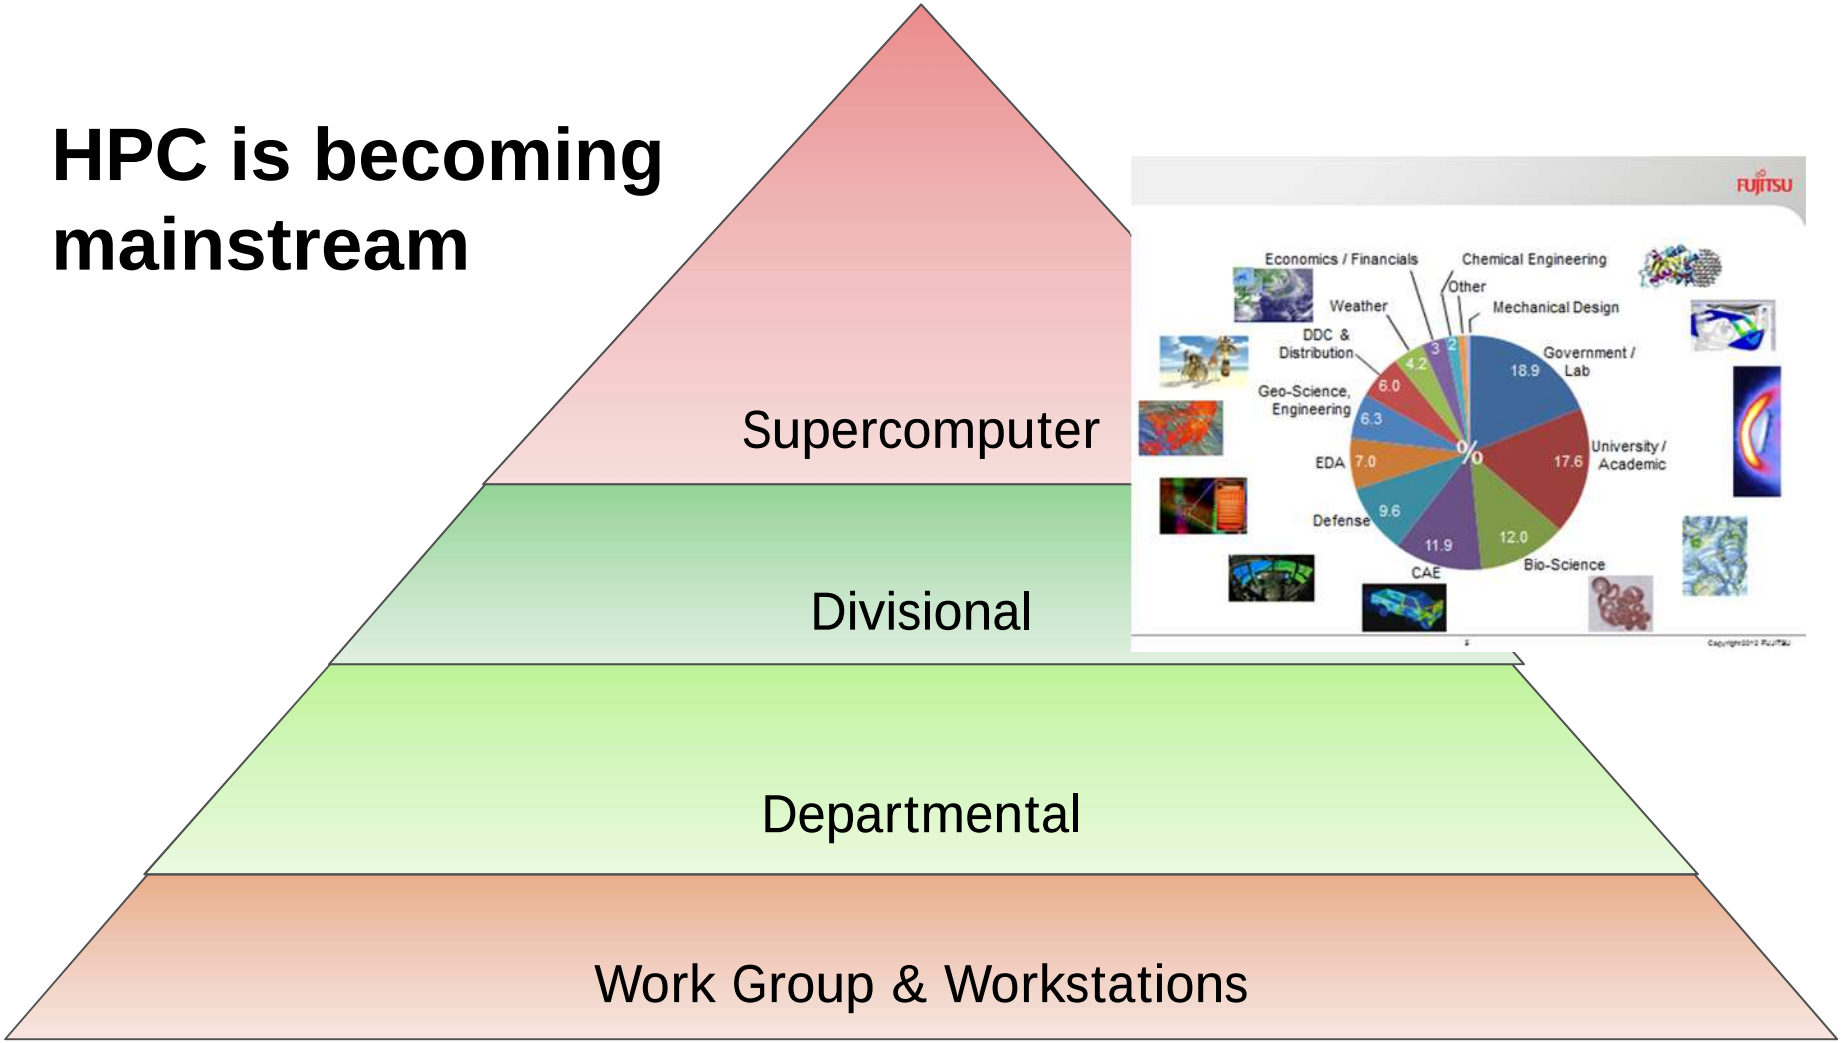

Fujitsu's approach to the HPC market

FUJITSU

Fujitsu covers the HPC customer needs with tailored HPC platforms

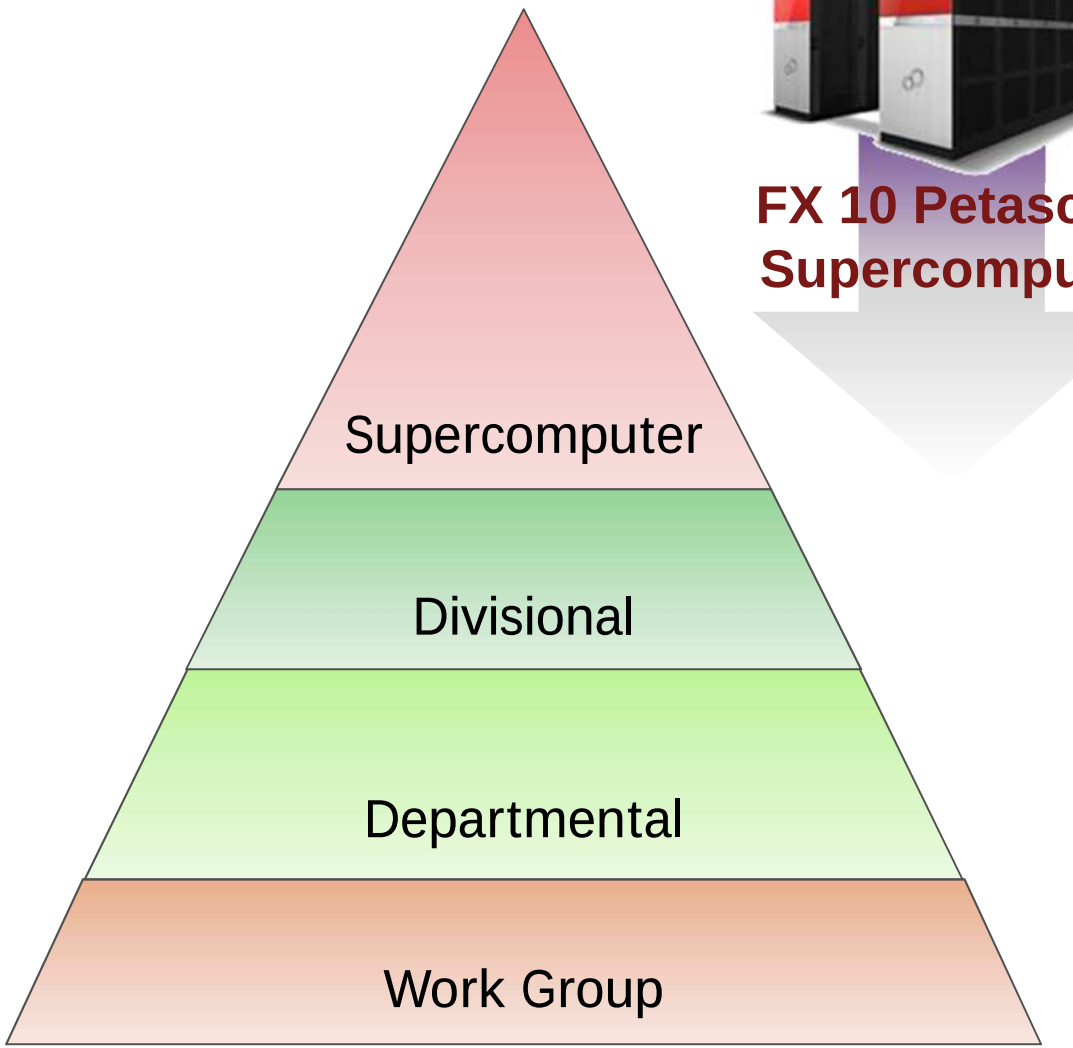

FX 10 Petascale Supercomputer

PRIMERGY HPC solutions

NEW

Power Workstation

FUJITSU

High Performance Computing with Fujitsu

HPC roots

K-Computer Development

Solution Positioning and Use

Fujitsu HPC Servers - past, present, future

FUJITSU

Fujitsu has been developing the most advanced supercomputers in the world since 1977!

PRIMERGY in HPC for 10 years!

No.1 in Top500 (June and Nov., 2011)

Exascale

Trans-Exascale

FX10

K computer

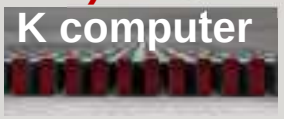

Vector Supercomputer Series

World's Fastest Vector Processor (1999)

VPP5000

NWT*

Developed with NAL

No.1 in Top500 (Nov. 1993)

Gordon Bell Pri (1994, 95, 96)

©JAXA

VPP300/700

VPP500

VP Series

F230-75APU

Japan's First Vector (Array) Supercomputer(1977)

Scalar MPP Series

AP3000

AP1000

SPARC Enterprise

FX1

Most Efficient Performance in Top500 (Nov. 2008)

PRIMEQUEST

PRIMEPOWER HPC2500

World's Most Scalable Supercomputer (2003)

Scalar Supercomputer Series

Next x86 generation

PRIMERGY CX400

PRIMERGY BX900 Cluster node

HX600 Cluster node

PRIMERGY RX200 Cluster node

Japan's Largest Cluster in Top500 (July 2004)

x86 Cluster Series

*NWT: Numerical Wind Tunnel

-1985 | 1990 | 1995 | 2000 | 2005 | 2010 | 2015 | 2020

"K computer" Achieves Goal of 10 Petaflops

FUJITSU

TOP 10 Systems - 11/2011		Cores	R_{max}	R_{peak}	Power
1	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	705024	10510.00	11280.38	12659.9
2	NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050	186368	2566.00	4701.00	4040.0
3	Cray XT5-HE Opteron 6-core 2.6 GHz				
4	Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU				
5	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows				
6	Cray XE6, Opteron 6136 8C 2.40GHz, Custom				
7	SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon 5570/5670 2.93 Ghz, Infiniband				
8	Cray XE6, Opteron 6172 12C 2.10GHz, Custom				
9	Bull bullx super-node S6010/S6030				
10	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband				

FUJITSU and RIKEN

- #1 on the TOP500 Lists June and November 2011
- #2 on the TOP500 List in June 2012
- LINPACK benchmark performance of 10.51 petaflops
- Received several big science application awards
- **Highest Performance Efficiency with 93%**

Fujitsu PRIMEHPC FX10

FUJITSU

Next level of supercomputing introduced Nov. ‘11

SPARC64™ VIII / IX fx RAS coverage ranges

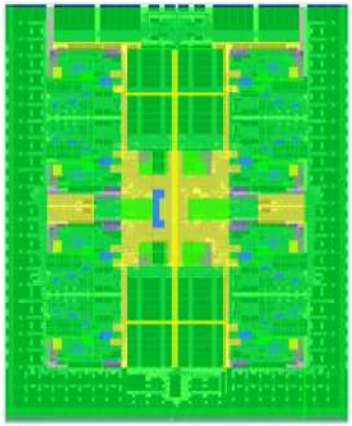

Hardware-based error detection + possible self-recovery range

Hardware-based error detection range

Range in which errors do not affect actual operation

Node	Theoretical computational performance	236.5 GFLOPS
	Processor	SPARC64™ IXfx (1.848GHz, 16 cores) x 1
	Memory capacity	32GB, 64GB
	Memory bandwidth	85GB/s
	Inter-node transfer rate	5GB/s × 2 (bidirectional) / link
System	No. racks	4 ～ 1,024
	Nodes	384 ～ 98,304
	Theoretical computational performance	90.80 TFLOPS ～ 23.24 PFLOPS
	Total memory	12 ～ 6,291 TB
	Interconnect	“Tofu” Interconnect
	Cooling Method	Direct water cooling + air cooling (Optional: Exhaust cooling unit)

FUJITSU

High Performance Computing with Fujitsu

PRIMERGY x86 based HPC

Modular HPC growth potential towards ….

FUJITSU

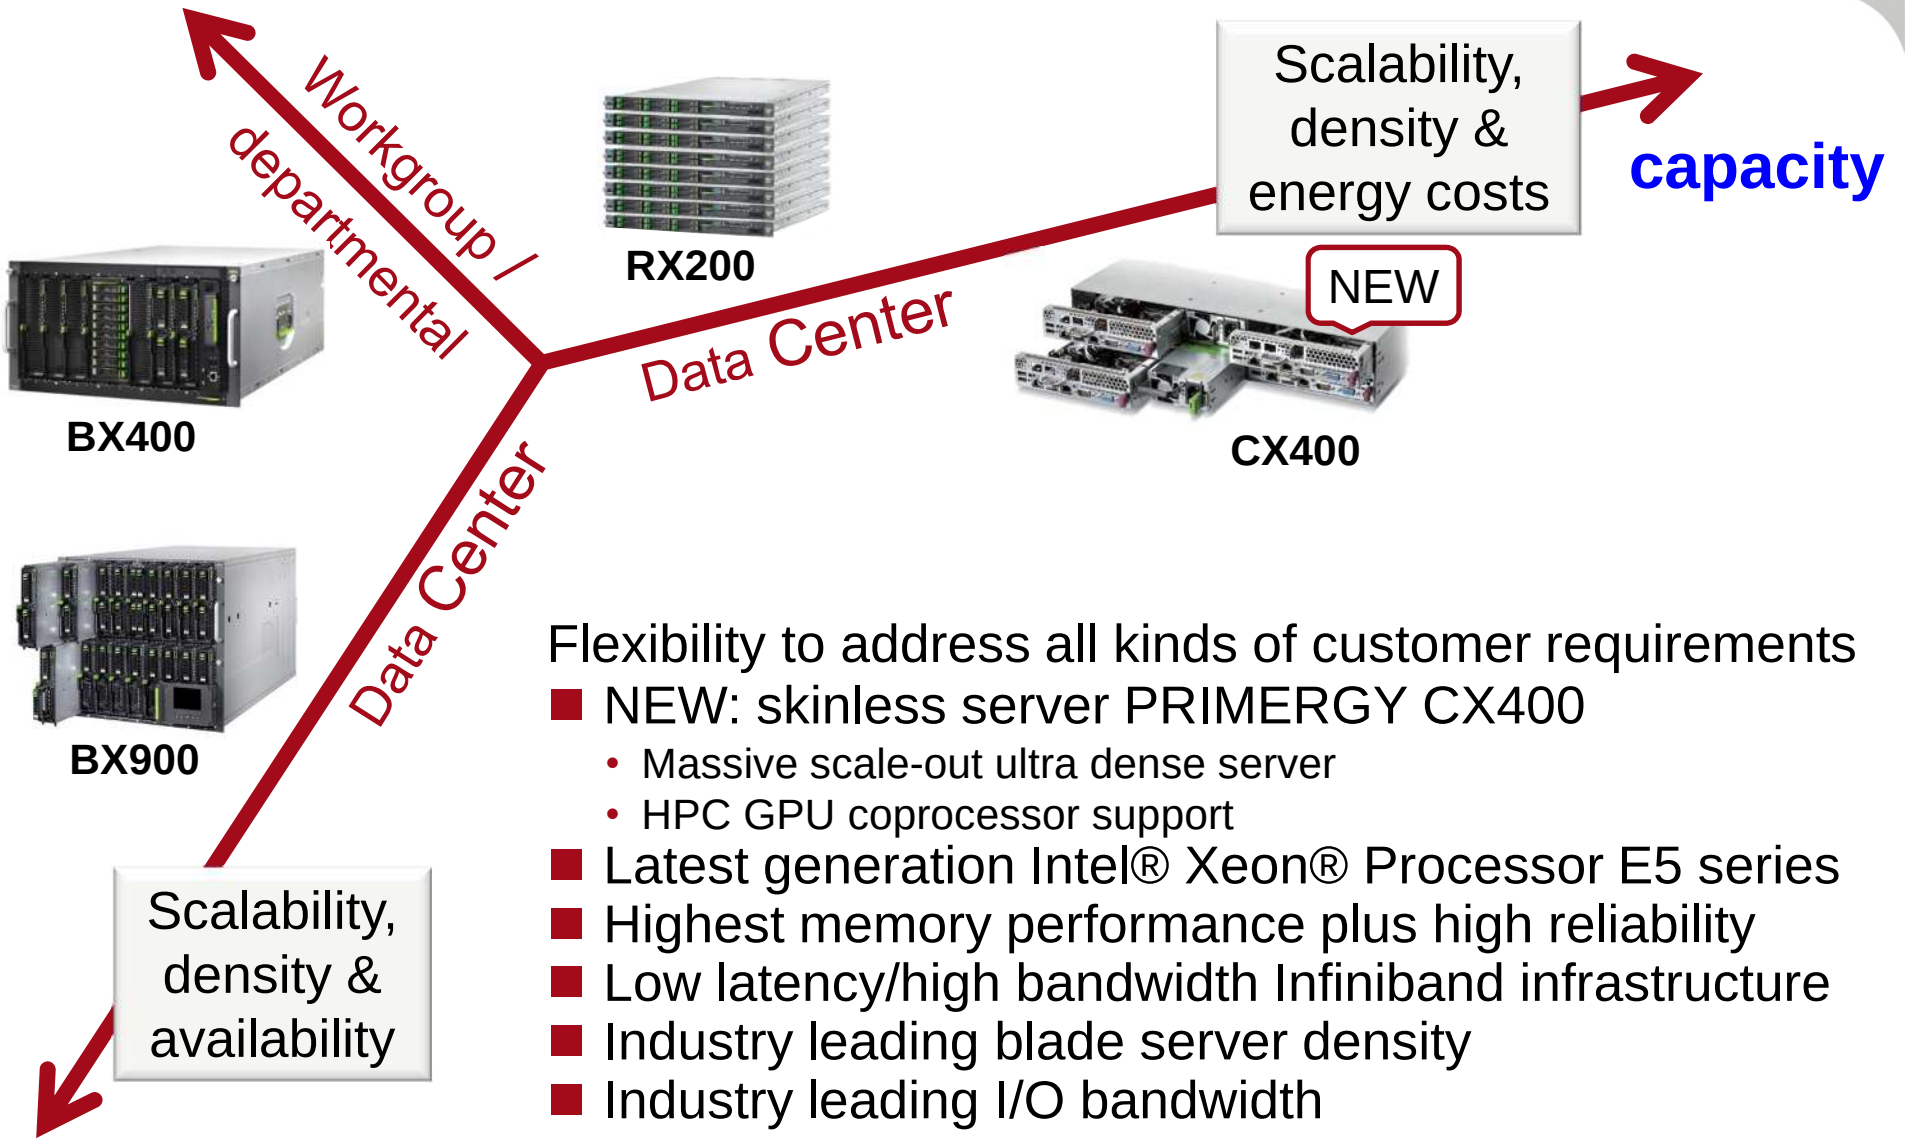

Flexibility to address all kinds of customer requirements

- NEW: skinless server PRIMERGY CX400
  - Massive scale-out ultra dense server
  - HPC GPU coprocessor support
- Latest generation Intel® Xeon® Processor E5 series
- Highest memory performance plus high reliability
- Low latency/high bandwidth Infiniband infrastructure
- Industry leading blade server density
- Industry leading I/O bandwidth

Most energy efficient server in the world

FUJITSU

- Fujitsu PRIMERGY achieves **world record in energy efficiency** and holds several best in class ratings
  - World record in SPECpower_ssj2008 by breaking the prestigious milestone of **6,000 overall ssj_ops/watt**
  - http://ts.fujitsu.com/ps2/press/read/news_details.aspx?id=6092

- Reduce energy consumption and current carbon footprint
  - Up to 73% more performance per Watt compared to the previous generation means
    - Up to 33% less energy for the same current performance level enables to better meet stringent environmental mandates for data centers
    - Up to 66% more workloads on current power budget without stressing current data center cooling

PRIMERGY CX400 S1: Scale-out Smart

Multi-Node System for HPC and Cloud Computing with compact granularity

- **PRIMERGY CX400 S1 doubles server density over traditional rack servers**
- New condensed „4in2U“ form factor
- More performance per U for large scale-out computing

4 servers, **each** with

- **two latest Intel Xeon E5-26xx** CPUs
- up to **512 GB Memory**
- up to **3 x** PCIe **Gen3** I/O slots

36 TB hot-plug storage drives

shared redundant power supplies and fans

„4in2U“

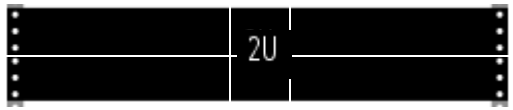

CX400 S1 - front

CX400 S1 - back

PRIMERGY CX2yy Server nodes at a glance

Double Performance per U /w condensed dual socket server nodes

- **CX250 S1 : 1U Server Node** hot-plug with 2 CPUs
  - standard node for HPC and Cloud Computing
  - 2 x latest Intel® Xeon® processor E5-2600 family
  - up to 512 GB RAM
  - 1 xPCIe expansion slot + 1 x mezzanine card

CX250 S1
half-wide 1U server node
4 x per CX400

- **CX270 S1: 2U Server Node** hot-plug with 2 CPUs + **GPGPU** option
  - HPC optimized node /w GPGPU acceleration
  - 2x Intel® Xeon® processor E5-2600 family
  - up to 512 GB RAM
  - 2 xPCIe expansion slots
  - 1 x GPGPU option: NVIDIA Tesla 20 series: M2075 / M2090

CX270 S1
half-wide 2U server node /w GPGPU option
2 x per CX400

2U height

- Up to 64 Intel Xeon processor cores
- up to 2.048 GB Memory
- up to 36 TB local storage

Building Blocks of PRIMERGY HPC Ecosystem

FUJITSU

PRIMERGY Server

Eternus Storage

Cluster Operation

PCM

FUJITSU

Edition

redhat.

Windows HPC Server 2008

ISV and Research Partnerships

ANSYS, MSC Software, SIMULIA, ESI GROUP, CD-adapco, LSTC, Altair, Exa, MAGMA

PreDiCT Initiative

Open Petascale Libraries Network

Consulting and Integration Services

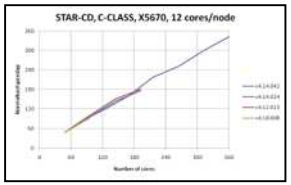

Sizing, design

Proof of concept

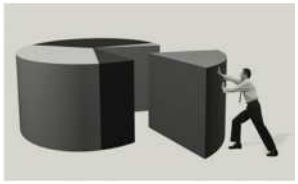

Integration into customer environment

intel Cluster Ready

Certified system and production environment

Complete assembly, pre-installation and quality assurance

Ready to operate delivery

Ready to Go

HPC Wales – A Grid of HPC Excellence

Initiative's motivation and background

- Positioning of Wales at the forefront of supercomputing
  - Promotion of research, technology and skills
- Improvement of economic development
  - **Creation of 400+ quality jobs, 10+ new business**

Implementation and rollout

- Distributed HPC clusters among 15 academic sites in Wales
  - Sophisticated tier model with central hubs, tier 1 and 2 sites, portal for transparent, easy use of resources
  - Final rollout of **PRIMERGY** x86 clusters in 2012

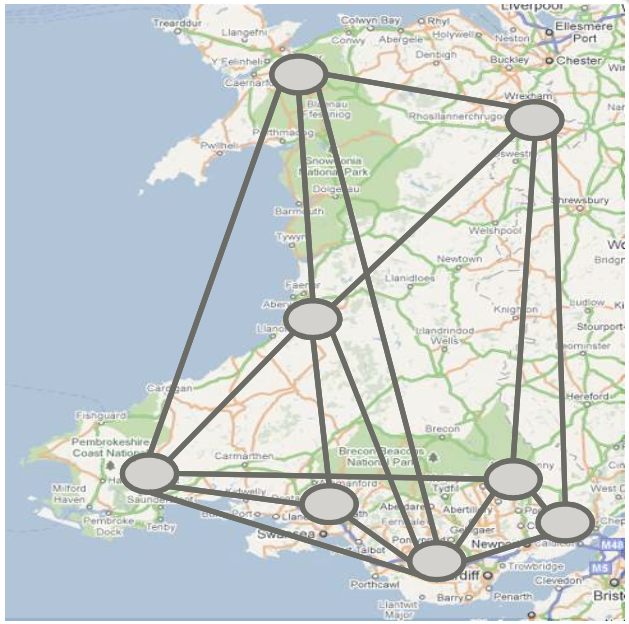

Performance & Technology

- **PRIMERGY CX250 and BX922**

- **~200 TFlops** aggregated peak performance
- **Infiniband**, 10 / 1 Gb Ethernet, FCS
- **Eternus** DX online SAN (home FS)
- **Parallel File System** (up to 10 GB/s)
  - DDN Lustre
- Backup & Archiving
  - Symantec, Quantum

Solution Design

- **User-focused solution** to access distributed HPC systems from desktop browser
- Multiple components integrated into a consistent environment with a **single sign-on**
- Data accessible across the entire infrastructure, **automated movement** driven by workflow
- **Collaborative sharing** of information and resources

Fujitsu Value Proposition

- **Best-in-class technology** combination
  - Latest Intel processor technology
  - Mix of Linux and Windows OS
  - Complete tuned software stack
- **Full service engagement**
  - Completely integrated design
  - Comprehensive engagement model at all levels (research, development, business)
- Professional **delivery management** and **governance**
  - End-to-end program management

Australia's National Computational Infrastructure

FUJITSU

Tokyo, June 15, 2012 — Fujitsu today announces its selection to deliver and install a new supercomputer, as well as to establish a Collaboration with the Australian National University's (ANU) National Computational Infrastructure (NCI) to provide high-end computational services to the Australian research community.

- x86 cluster-type supercomputer based on Fujitsu **PRIMERGY** servers
- the largest PRIMERGY deployment worldwide
- Theoretical performance of **1.2 petaflops**
- one of the world's fastest supercomputers
- deployed in a new data center in Canberra
- starting in mid-December 2012 and handover in early 2013.

FUJITSU

High Performance Computing with Fujitsu

Future Technologies and Trends

- Parallelism next Level
- HPC in the Cloud

The Future of HPC

FUJITSU

Intel Xeon Processor
1 core, 2 threads

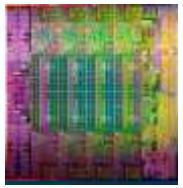

Intel Xeon E5 Processor
8 cores,16 threads

Towards
Many Core
Architectures

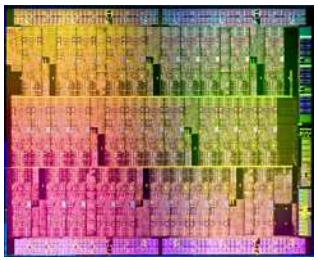

A path towards Exascale enforces a deployment of parallelism at each level to the ultimate possible extent

- Node level (distributed memory)
- Multi socket (shared memory on nodes)
- CPU level (number of cores)
- Instruction level (SIMD)

Intel
Knights Ferry,
Knights Corner

Nvidia
Tesla GPGPU
(M-Series)

Challenges

- Increasing parallelism within the CPU results in demand for higher memory bandwidth and thus greater complexity of the memory hierarchy
- Node parallelism enforce the development and deployment of ultra-high-speed interconnect
- Increasing parallelism towards millions of cores leads to increase in system errors
- **Amdahl's Law** is more alive than ever and demonstrates that even the smallest portion of serial code dominates (negatively) the overall performance of a code
  - A tiny portion of just 1 per mill of sequential code within a parallel application strictly limits the possible speedup to a factor of 1000

HPC in the Cloud

FUJITSU

- **HPC in the Cloud – a different HPC delivery model**
- **Benefits**: cost savings (capex vs opex); biz flexibility (resources when you need them); scale up & down (as much as you need)
- **Challenges** for HPC Cloud have to be overcome:
  - **Security** and privacy guarantees (to protect corporate assets like software, data, and personal information)
  - **Performance** degradation: data transfer, virtualization, sub-optimal h/w
  - **ROI** is often not clear: in-house HPC versus Cloud HPC
  - **Learning** curve may take time and money
  - **Mental** challenge: loosing control: hardware, software, data => TRUST
- **Fujitsu Technical Computing Cloud in Aug 2011 in Japan; early users are running test and production jobs**

Fujitsu Technical Computing Cloud

FUJITSU

- **Started** in August 2011 with 3 services for technical computing in manufacturing:
  - **Analytical Platform Service**, which provides a platform for conducting analytical simulations;
  - **Analytical Applications Service**, which provides analytical applications; and
  - **Analytical Help Desk**, which provides support for setting up and running analytical simulations.
- **Use Case**: Nikon Corporation, a manufacturer of precision optical equipment:
  - First, using the Analytical Platform Service to cope with the electricity supply shortages in Japan in 2011
  - Then, running analytical simulations for designing and developing optical equipment such as cameras, photolithography machines, and measuring instruments.

Fujitsu Global HPC Office in Munich

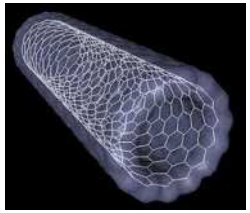

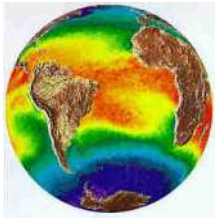

- **FUJITSU has opened Global HPC Office in Munich, 1st April 2012**
- **HPC business promotion for global HPC market**
- **HPC Go To Market operation with enriching HPC offering**

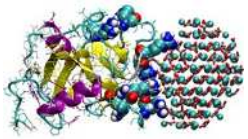

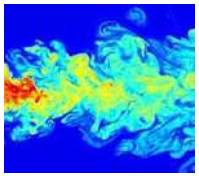

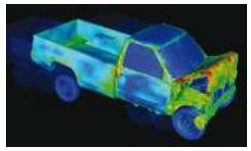

Fujitsu Global HPC Office in Munich

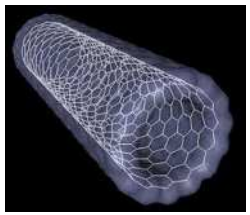

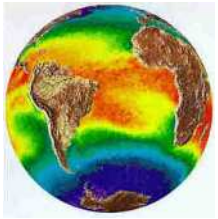

Global
HPC Office

- **FUJITSU has opened Global HPC Office in Munich, 1st April 2012**
- **HPC business promotion for global HPC market**
- **HPC Go To Market operation with enriching HPC offering**

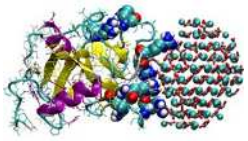

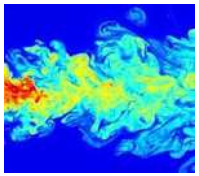

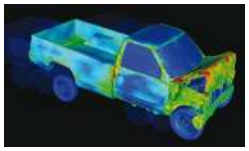

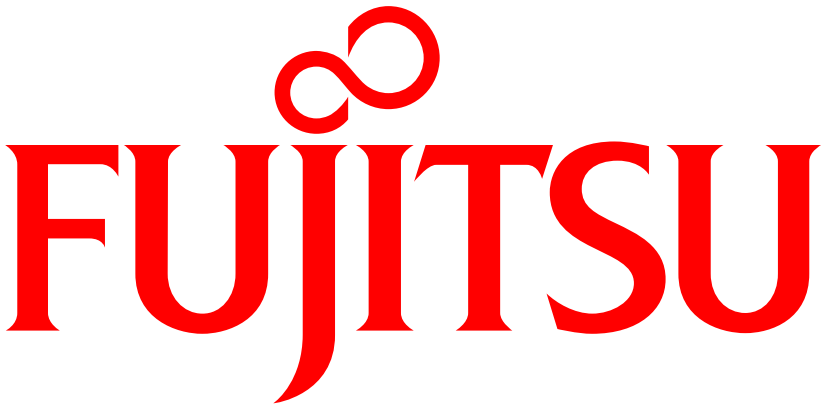

shaping tomorrow with you