FUJITSU

shaping tomorrow with you

Ready-to-go with Ready-to-use

# PRIMERGY x86 – HPC Simplicity

# The Concept of HPC Simplicity

FUJITSU

## Ready-to-go HPC Platform

Validated systems
- Optimal configuration
- Certified operation
- HPC software stack

## PRIMERGY HPC Portal

Adapted for HPC
- Application on-boarding
- Pre-defined workflow catalogue
- Collaboration and teamwork

Built-in expertise

## HPC Ready-to-use Solution

# *Ready-to-go* systems for HPC

- PRIMERGY x86 server configuration and delivery

# Building Blocks of PRIMERGY HPC Ecosystem

FUJITSU


**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**

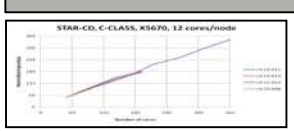

Sizing, design

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**

# Modular HPC growth potential towards ….

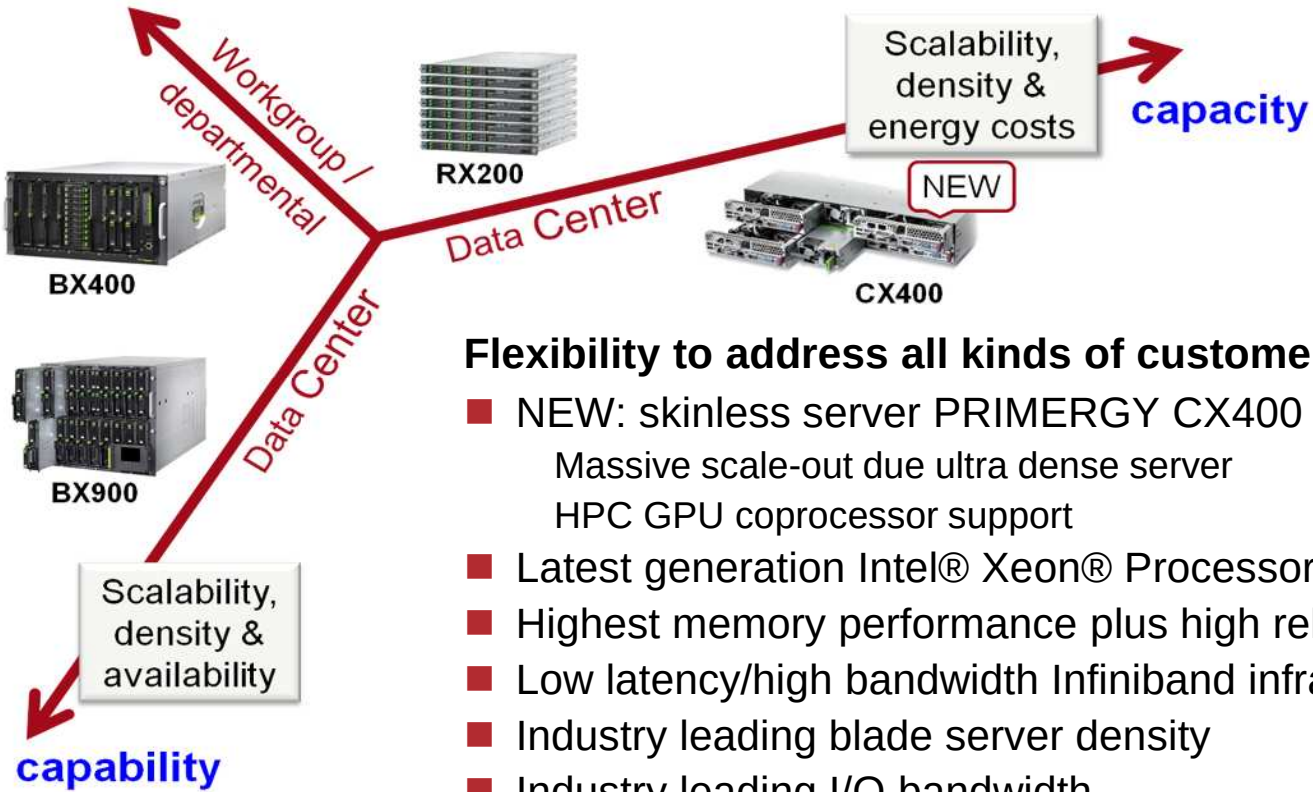


**Flexibility to address all kinds of customer requirements**

- NEW: skinless server PRIMERGY CX400
  - Massive scale-out due 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 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 - HPC Design

FUJITSU

CX400 combines high performance with high density at lower overall investment

## High Density / Scalability in 2U Chassis

- HPC requirements optimally fulfilled
  - Up to 4 nodes (1U) or 2 nodes (2U) per 2U chassis
  - 2x Intel® Xeon® **E5-2600** processors / node
  - Intel® Xeon® processor E5-2400 node coming soon
  - 16 DIMMs, up to **1600MHz**
  - Redundant, hot-plug PSUs for enhanced availability / lower servicing effort
  - Up to 24x HDD
  - **FDR Infiniband** interconnect option for highest, most efficient bandwidth and lowest latency
  - **GPU** Option (2U node)
  - Support of Intel MIC Q1/ 2013 planned

## Main usage scenario

**Cloud** **HPC**

# The Future in High Performance Computing

FUJITSU

Intel Xeon Processor
1 core, 2 threads

Intel Xeon E5 Processor
8 cores,16 threads

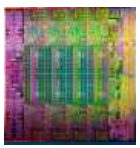

Towards Many Core Architectures

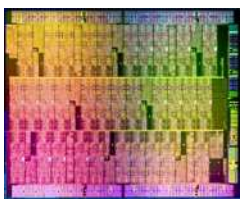

A path towards Exascale enforces a **deployment of parallelism** at each level to the ultimate possible extent

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

## 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

# Many Core Architectures

FUJITSU

CX400/CX270 and Floating Point accelerators address the core level parallelism

Fujitsu CX400

Intel Knights Ferry, Knights Corner


NvidiaTesla GPGPU (M-Series)

| | Intel Sandy Bridge E5-2600 | Intel MIC Knights Corner | Nvidia Fermi GPU Tesla 2050 |
|---|---|---|---|
| **Availability** | Today | Q4/2012 | Today (Kepler in Q4/2012) |
| **Cores** | 8 (superscalar) | > 50 | Tesla 2050: 448 |
| **Programming** | Standard languages | Standard languages and OpenCL | CUDA, OpenCL |
| **Architecture*** | Multicore | Manycore standalone (possibly) & hybrid | Hybrid |

# *Ready-to-use* PRIMERGY HPC solutions

- PRIMERGY HPC Portal with expertise in-built

# PRIMERGY HPC Portal total solution

FUJITSU

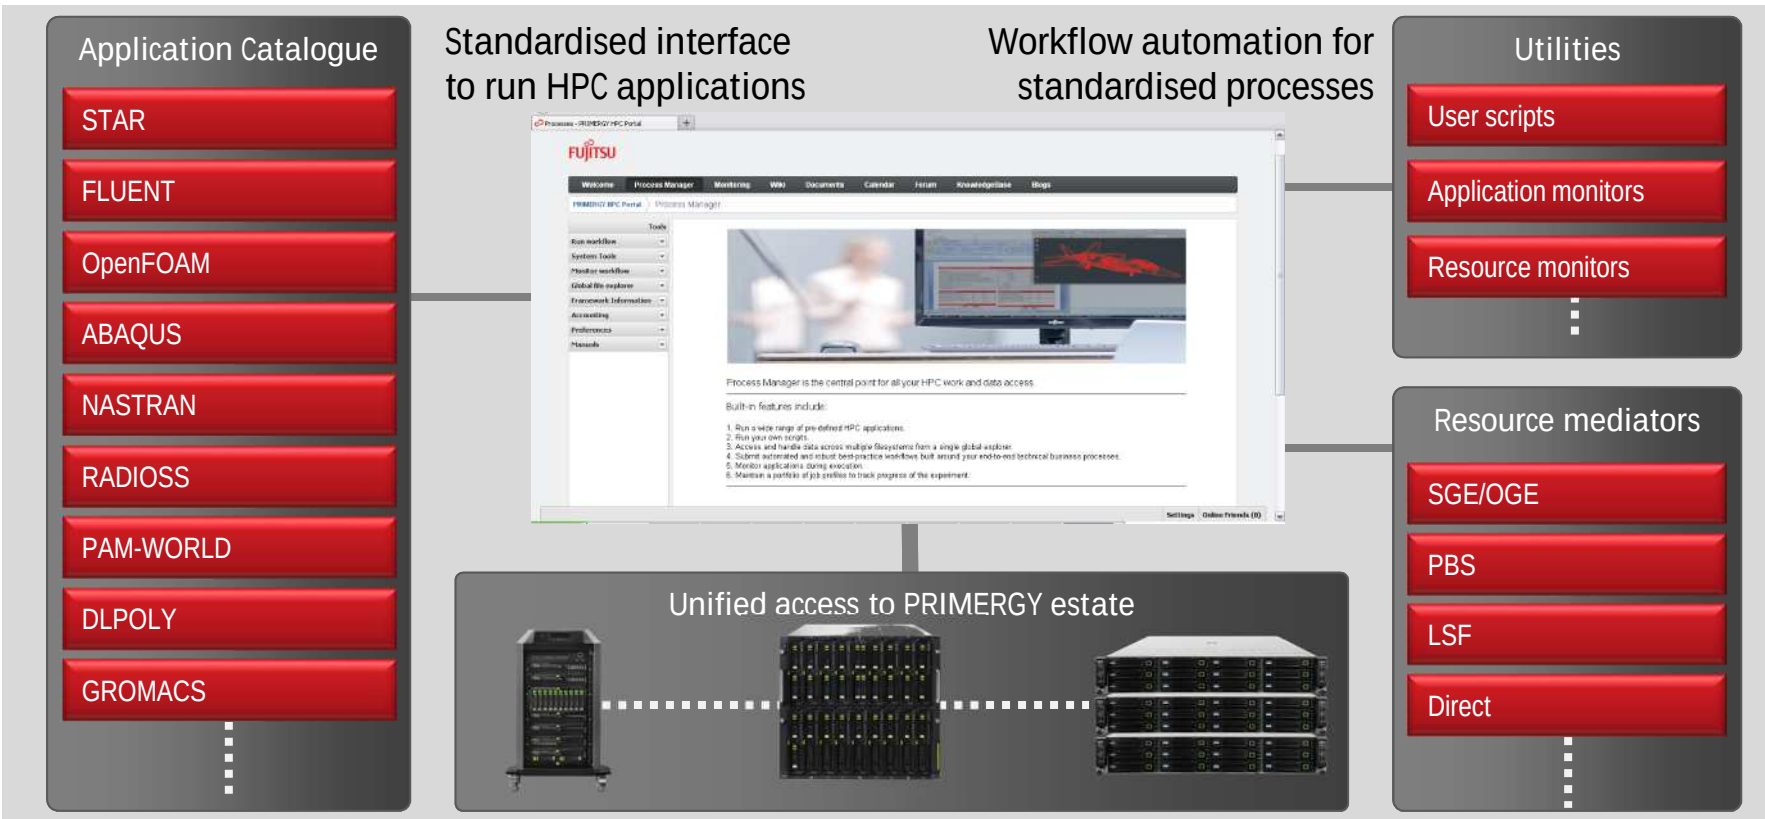

# Running an HPC Application

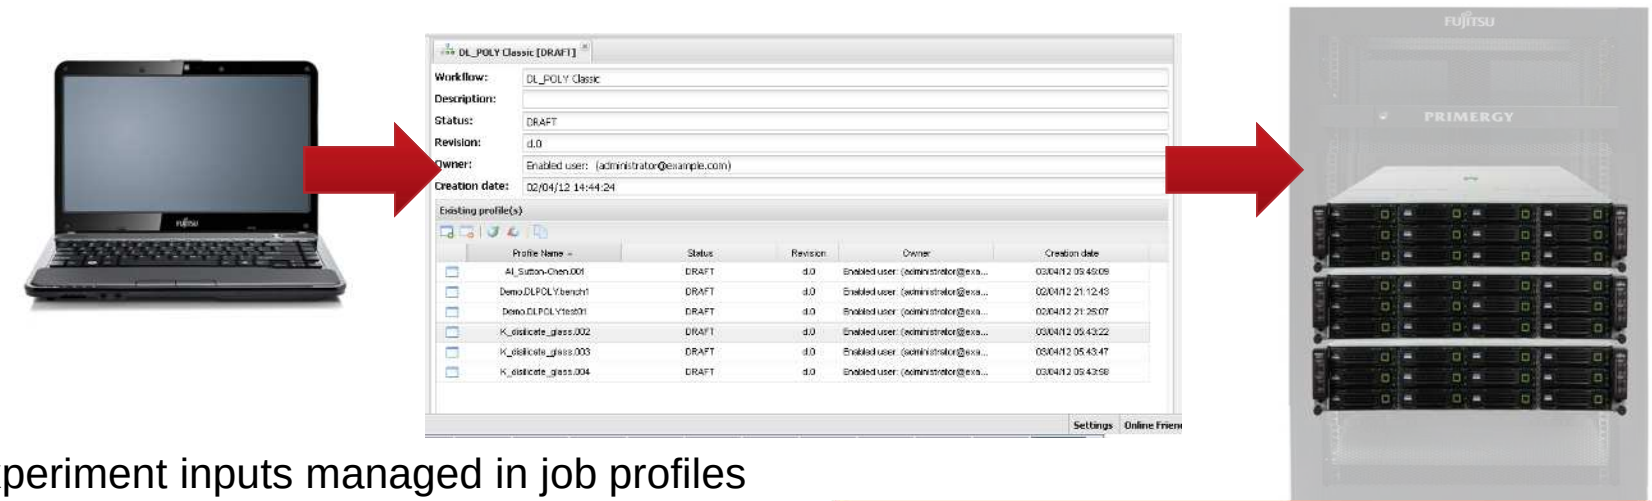


- Experiment inputs managed in job profiles
  - Systematic traceable approach over experiment lifecycle
  - Shareable – project team, between service designer and end-users
  - Reduced operations to submit

Usable with minimal learning
By end-users with little/no IT knowledge
For faster, more accurate setup

# HPC Simplicity – In-built Expertise

FUJITSU

## HPC Simplicity

- Make HPC easier to use
- Share and exchange more widely
- Use resources more effectively
- Broaden HPC access and process reuse

Point-click parameterisation and global file access

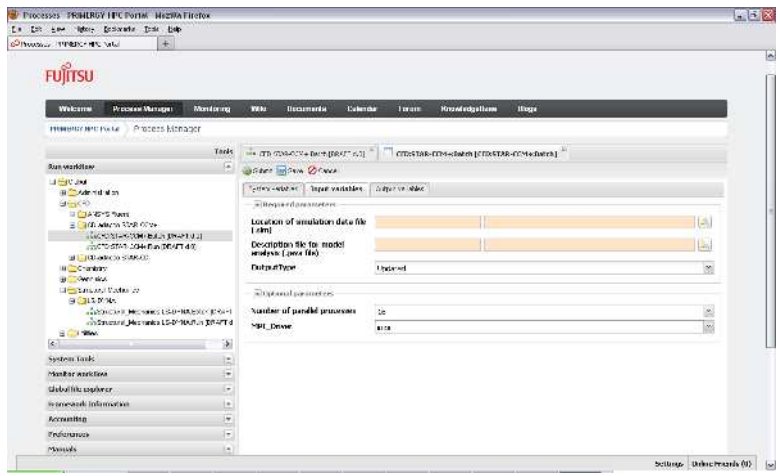

# HPC Simplicity – In-built Expertise

FUJITSU

## HPC Simplicity

- Make HPC easier to use
- Share and exchange more widely
- Use resources more effectively
- Broaden HPC access and process reuse

Team knowledge and project management

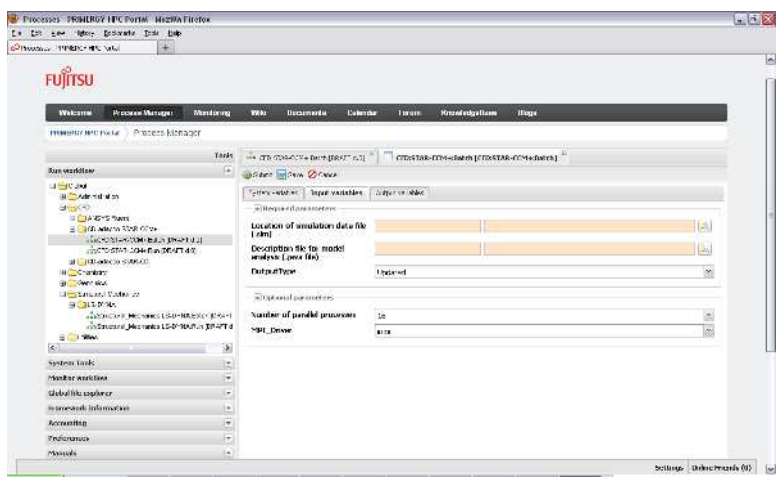

# HPC Simplicity – In-built Expertise

FUJITSU

## HPC Simplicity

- Make HPC easier to use
- Share and exchange more widely
- Use resources more effectively
- Broaden HPC access and process reuse

Project and application usage reporting

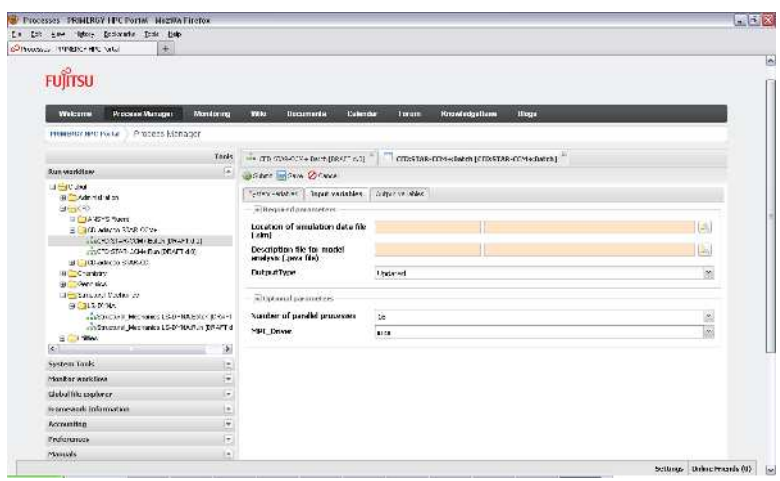

# HPC Simplicity – In-built Expertise

FUJITSU

## HPC Simplicity

- Make HPC easier to use
- Share and exchange more widely
- Use resources more effectively
- Broaden HPC access and process reuse

Catalogue of HPC application workflows

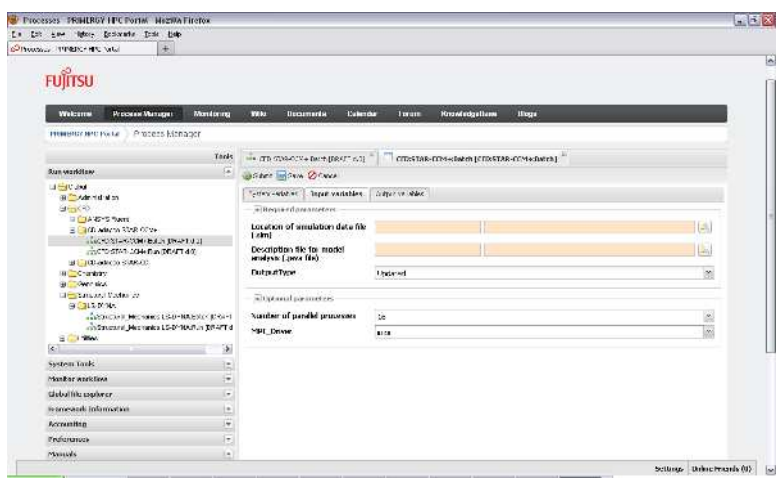

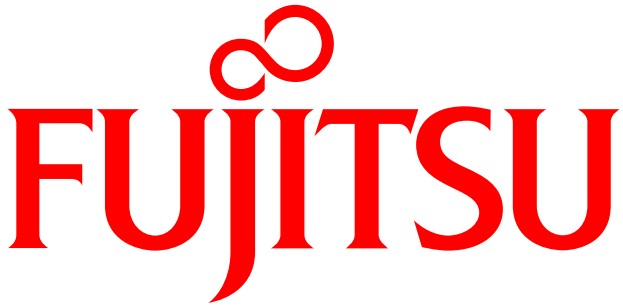


shaping tomorrow with you