

# Fujitsu Petascale Supercomputer PRIMEHPC FX10



4x2 racks (768 compute nodes) configuration

# PRIMEHPC FX10 Highlights




- Scales up to 23.2 PFLOPS
- Improves **Fujitsu's supercomputer technology** employed in the FX1 and the "K computer," the world's fastest supercomputer.

Hybrid parallel  
(VISIMPACT)  
Collective Sw  
FX1



SPARC64™ VII 4C/40GF  
121TFLOPS, CY2008 ~

Tofu interconnect  
ISA extension (HPC-ACE)  
K computer



SPARC64™ VIIIfx 8C/128GF  
11PFLOPS, CY2010 ~

## PRIMEHPC FX10



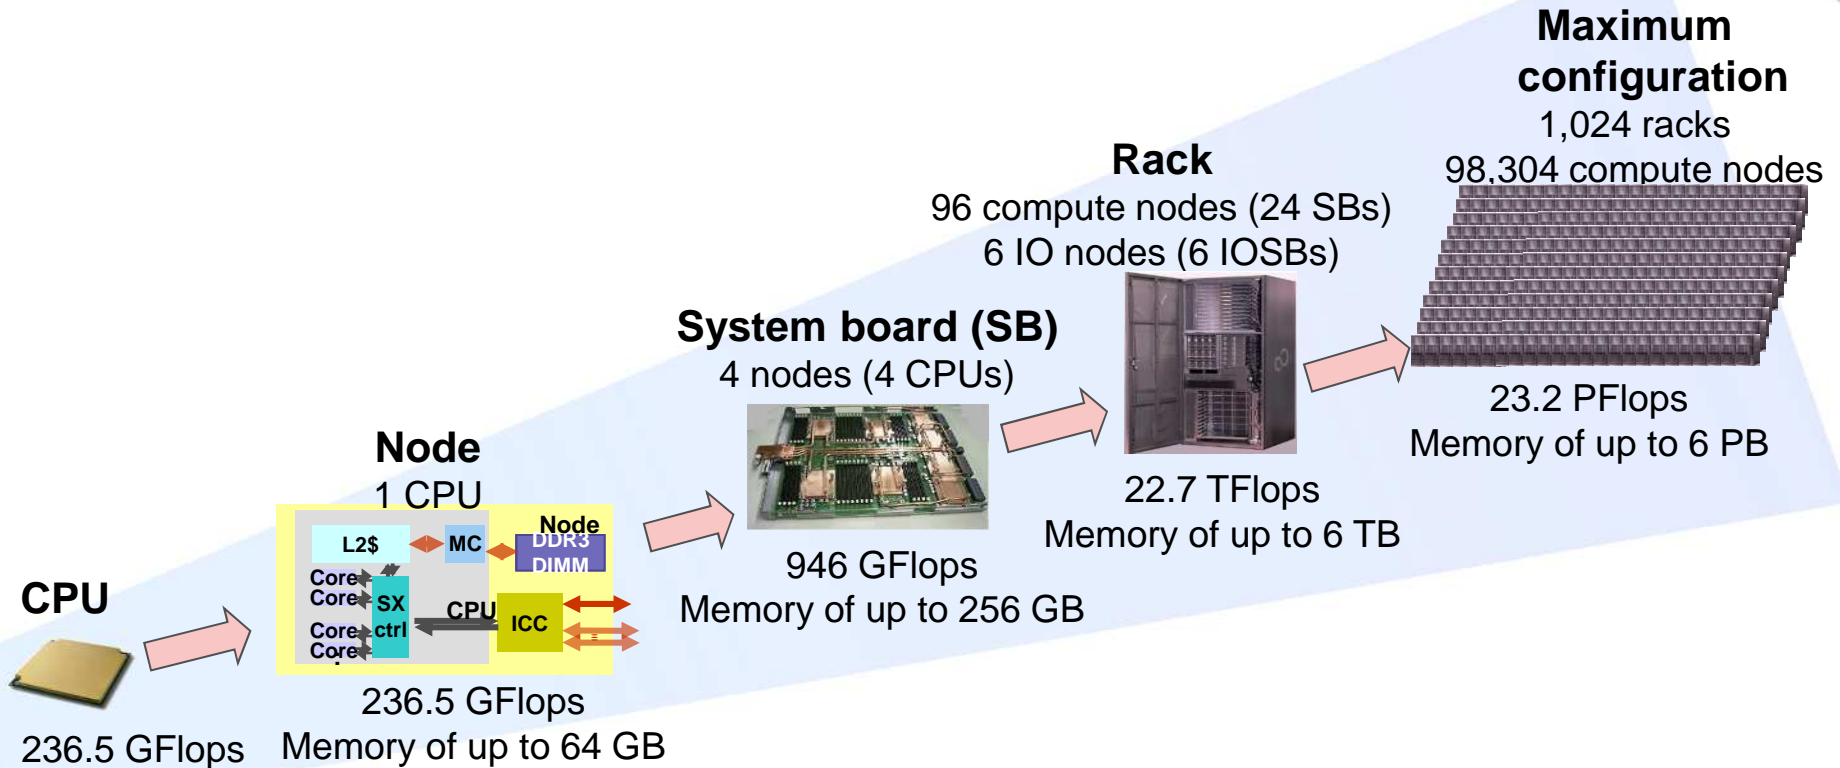
SPARC64™ IXfx 16C/236.5GF  
~ 23PFLOPS, CY2012 ~

# PRIMEHPC FX10 Features



- **High-speed and ultra-large-scale computing environment**
  - Up to 23.2 PFLOPS (98,304 nodes, 1,024 racks, 6 petabytes of memory)
- **SPARC64™ IXfx – high performance w/ low power consumption**
  - 236.5 GFLOPS, and over 2GFLOPS/W
- **High execution performance with massively parallel apps**
  - Tofu interconnect, Technical Computing Suite, and VISIMPACT
- **High reliability and high operability for a large-scale system**

# Implementation of PRIMEHPC FX10



# SPARC64™ IXfx




## ■ Inherited features from SPARC64™ VIIIfx

- High performance
  - HPC-ACE
  - Embedded memory controller
  - Hardware barrier mechanism
- Low power
  - Moderate frequency
  - Dynamic/Leakage power reduction
- High reliability
  - RAS features

## ■ New features

- Doubled number of cores
- Doubled size of L2\$



SPARC64™ IXfx


# SPARC64™ VIIIfx / IXfx Specifications



		SPARC64™ VIIIfx	SPARC64™ IXfx
<b>Number of cores</b>		8 cores	16 cores
<b>Clock frequency</b>		2 GHz	1.848 GHz
<b>Cache</b>	L1	I: 32KB/core, D: 32KB/core	I: 32KB/core, D: 32KB/core
	L2	6MB (Shared cache)	12MB (Shared cache)
<b>Peak performance</b>		128 Gigaflops	236.5 Gigaflops
<b>Memory throughput</b>		64 GB/s	85 GB/s
<b>Power consumption</b>		58W	110W
<b>Process</b>		45 nm	40 nm
<b>Die size</b>		22.7 mm × 22.6 mm	21.9 mm × 22.9 mm
<b>Number of transistors</b>		Approximately 760 million	Approximately 1.87 billion

# Tofu Interconnect

- Highly scalable and usable direct network (6D mesh/torus)
  - 10 redundant high BW links, 4 RDMA engines (4x2 simultaneous transfer)
  - Good collective communication performance with Tofu original algorithms
- Tofu barrier for barrier & reduction in H/W
- Direct attached interconnect controller



**Tofu realizes scalable systems beyond 100,000 nodes  
With low power consumption, low latency, and high BW**

# PRIMEHPC FX10 Software Structure



## User/ISV Applications

## HPC Portal / System Management Portal

### System operations management

- System configuration management
- System control
- System monitoring
- System installation & operation

### High-performance file system

- Lustre-based distributed file system
- High scalability
- IO bandwidth guarantee
- High reliability & availability

### Compilers

- Hybrid parallel programming
- Sector cache support
- SIMD / Register file extensions

### Support Tools

- IDE
- Profiler & Tuning tools
- Interactive debugger

### Job operations management

- Job manager
- Job scheduler
- Resource management
- Parallel execution environment

### VISIMPACT™

- Shared L2 cache on a chip
- Hardware intra-processor synchronization

### MPI Library

- Scalability of High-Func.
- Barrier Comm.

File system, operations management

Application development environment

## Linux-based enhanced Operating System

- Enhanced hardware support
- System noise reduction
- Error detection / Low power

## PRIMEHPC FX10



# PRIMEHPC FX10 Specifications




System specifications	Number of racks	4 to 1,024
	Number of compute nodes	384 to 98,304
	Theoretical Peak performance	90.8 Tera FLOPS to 23.2 Peta FLOPS
	Total memory capacity	12 TB to 6 PB
	Interconnect	Tofu interconnect
	Cooling method	Direct water cooling + air cooling Option : Exhaust cooling unit
Node specifications	Theoretical peak performance	236.5 Giga FLOPS
	Processor	SPARC64 IXfx (1.848 GHz, 16 cores) x 1
	Memory capacity	32 GB, 64 GB
	Memory bandwidth	85 GB/s
	Interconnect link bandwidth	5 GB/s x 2 (bi-directional)

# PRIMEHPC FX10 Summary



- Improves Fujitsu's supercomputer technology employed in the "K computer," the world's fastest supercomputer
  - Newly developed SPARC64 IXfx processor (236.5GFLOPS)
  - Tofu interconnect (scales up to 98,304 nodes)
- Fujitsu provides all hardware and software
  - Processor and interconnect
  - Technical Computing Suite HPC middleware, as well as FEFS high-performance distributed file system



shaping tomorrow with you