FUJITSU

shaping tomorrow with you

Key Hardware Technologies for the Next-Generation PRIMEHPC – Post-FX10

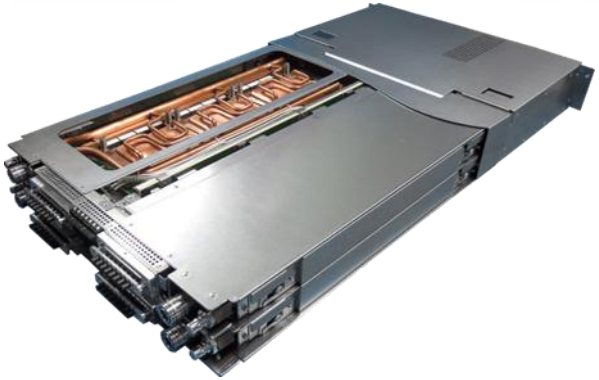

Key hardware technologies for Post-FX10

FUJITSU

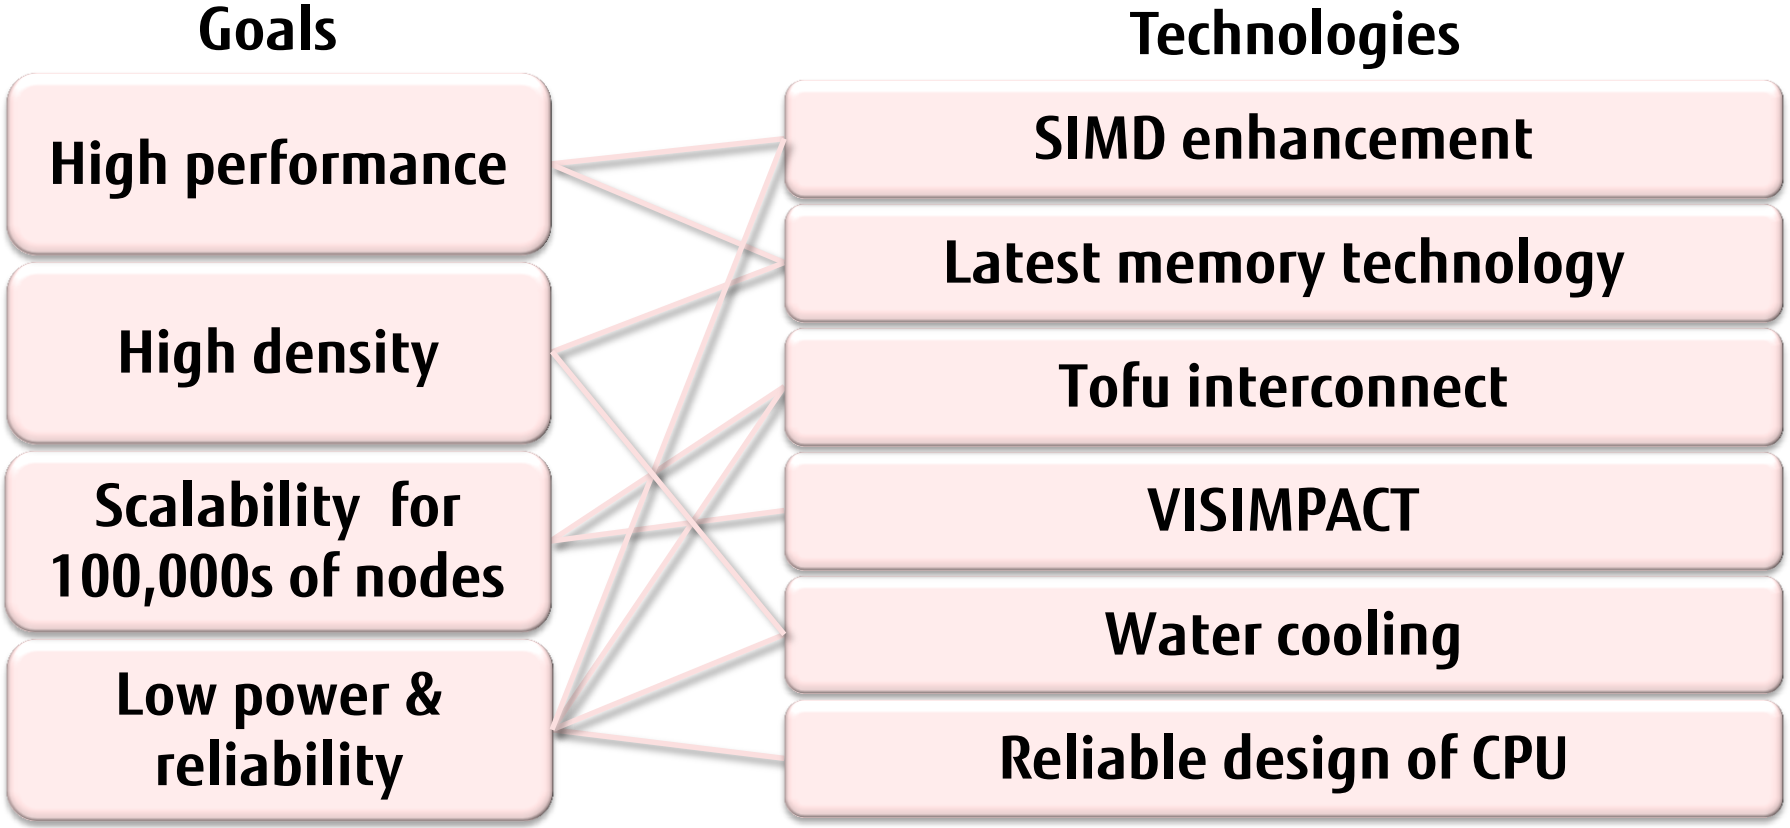

Key hardware technologies for Post-FX10

FUJITSU

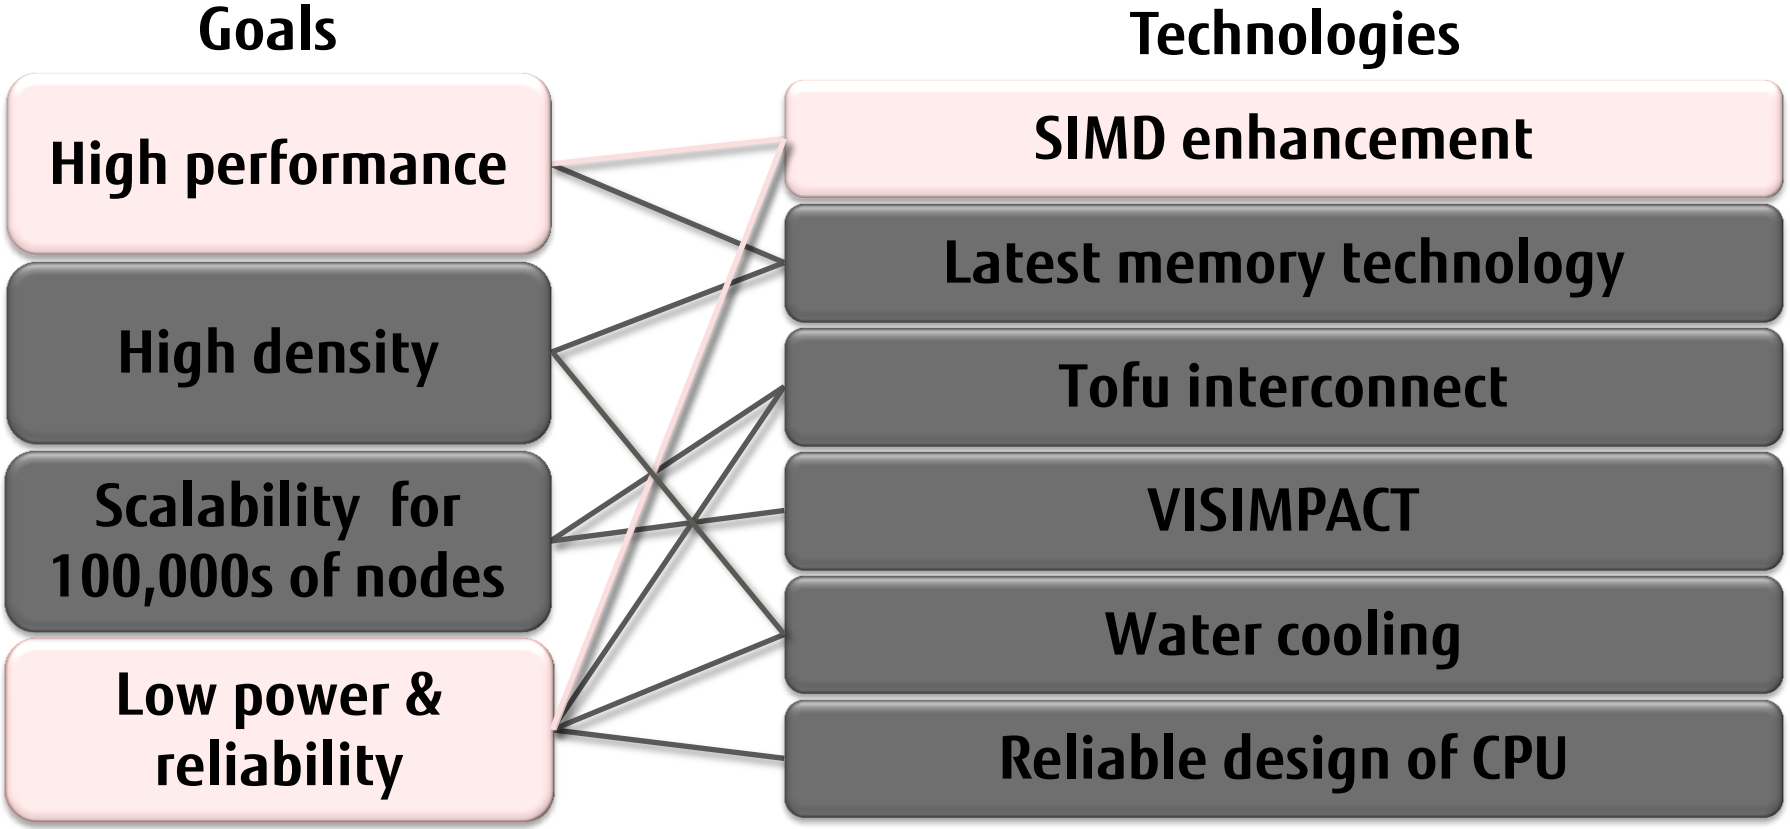

SIMD enhancement of Post-FX10

FUJITSU

- Wider SIMD
  - Various functions for real application performance
  - Increases single precision performance 2x over double precision calculations
  - 8-byte integer SIMD

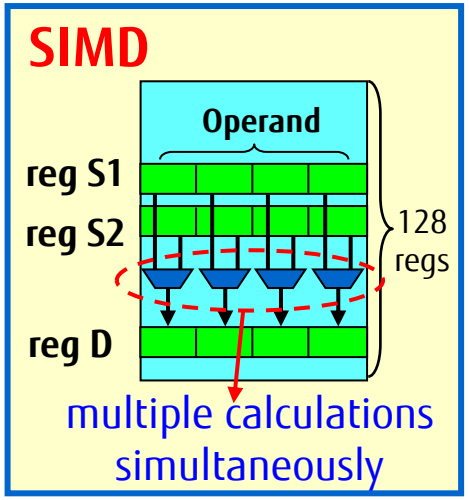

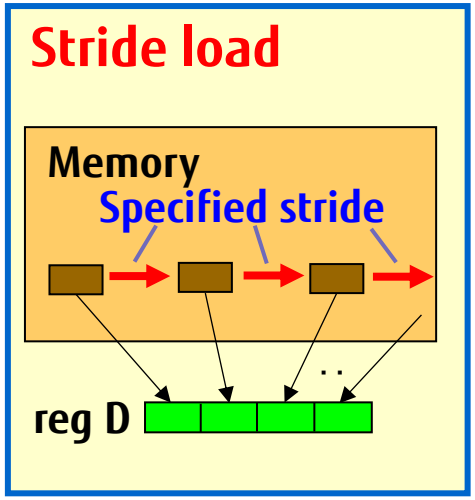

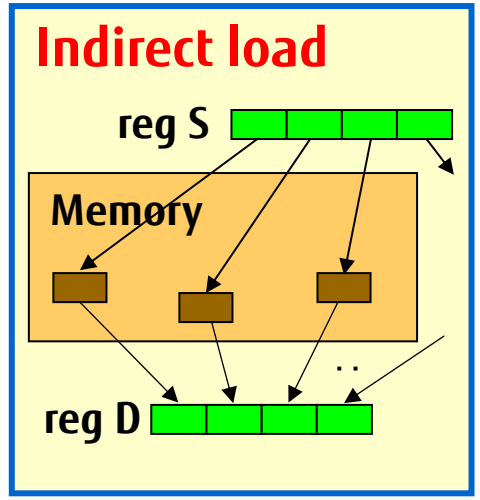

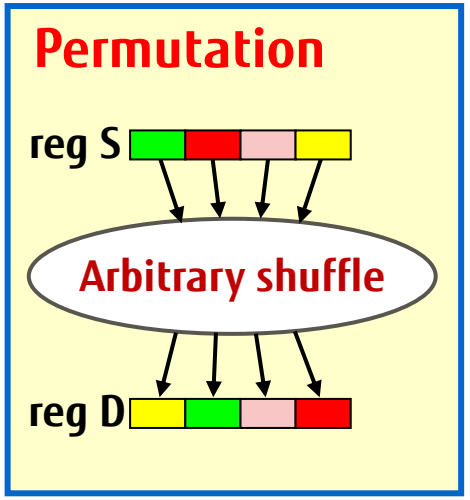

Key hardware technologies for Post-FX10

FUJITSU

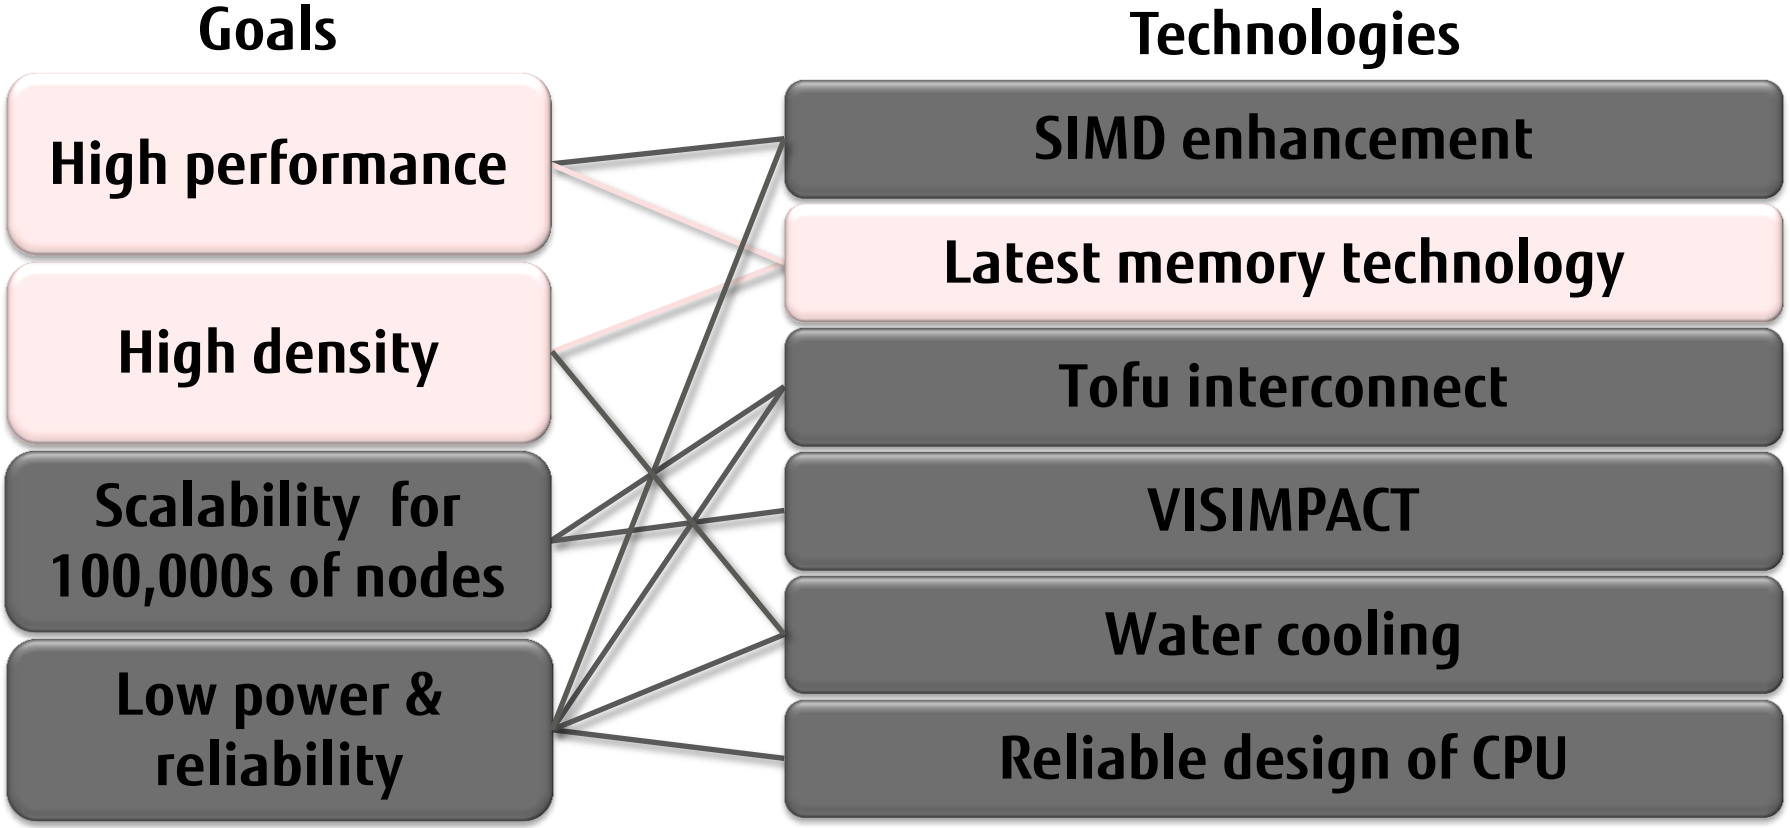

Latest memory technology, HMC

FUJITSU

- Hybrid Memory Cube
  - High bandwidth for application performance
  - High capacity for higher density

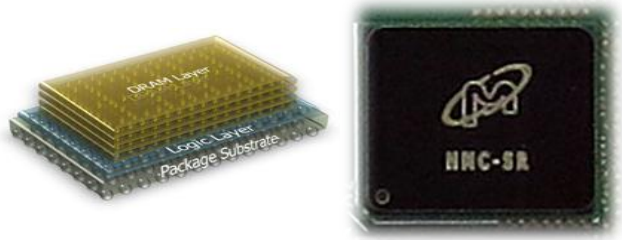

	Capacity/package	Bandwidth/package	Other concern
HMC x8	Good	Very good	
HBM* x8	Fair	Very good	Cost/SCM of 2.5D
DDR4-DIMM x8	Very good	Low	
GDDR5 x8	Low	Good	No successor

*HBM: High Bandwidth Memory

Key hardware technologies for Post-FX10

FUJITSU

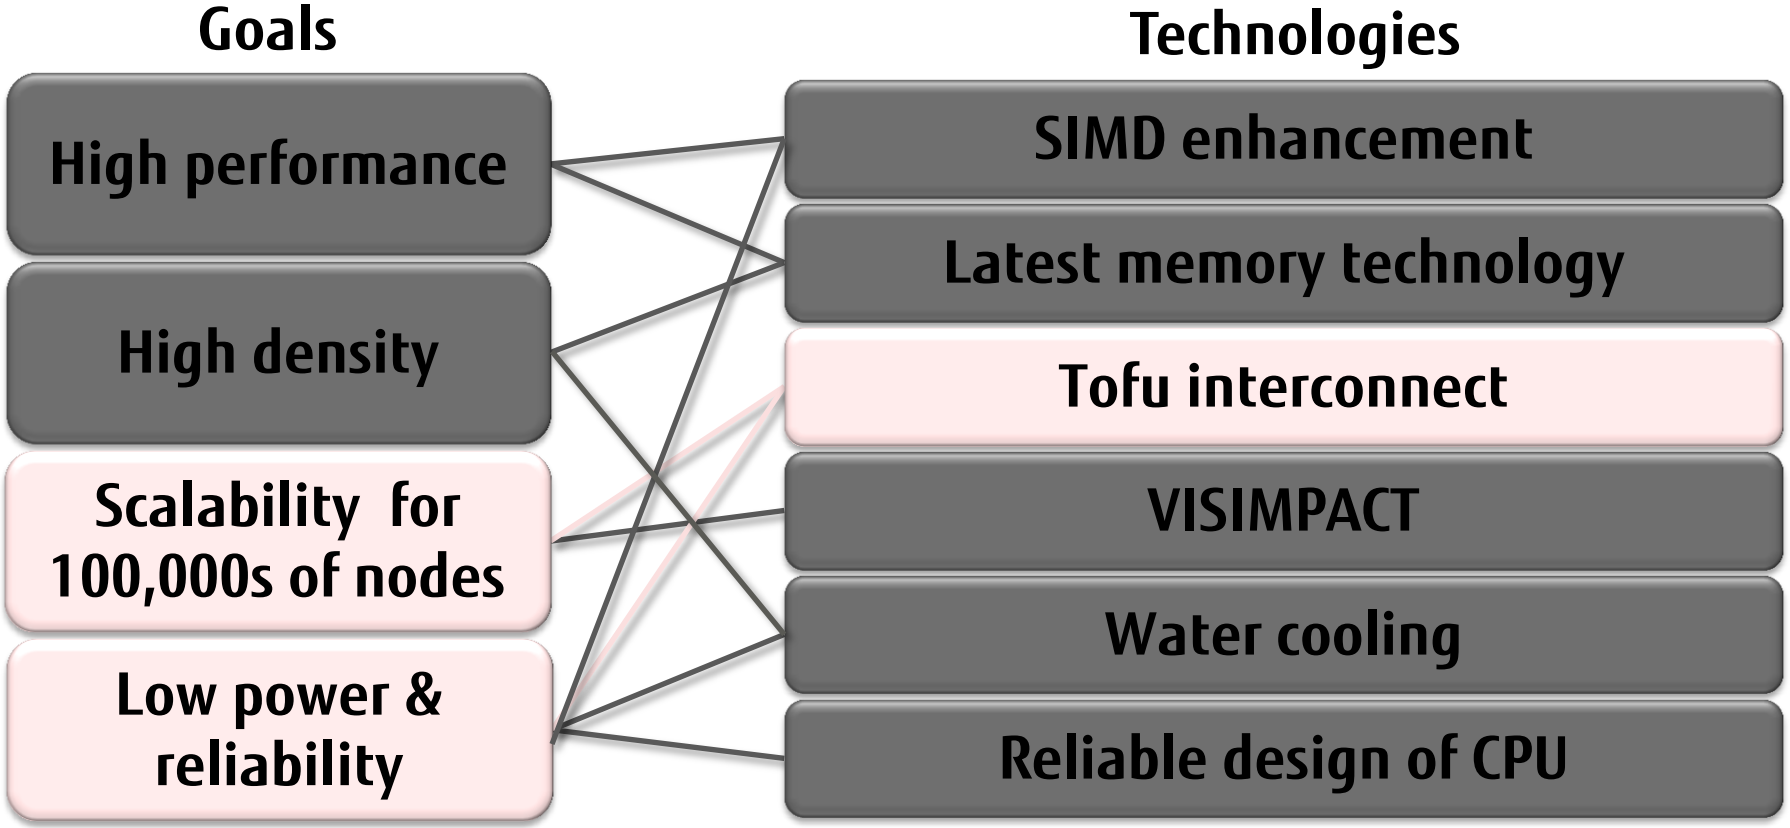

Tofu interconnect

FUJITSU

- Scalable beyond 100,000 nodes
  - 6-Dimension mesh/torus direct network

    Low average hops and high bisectional bandwidth
  - High operability by using redundant connections
  - Hardware collective communication support
- Tofu2 for Post-FX10
  - Bandwidth and latency optimized
  - Optical connection support

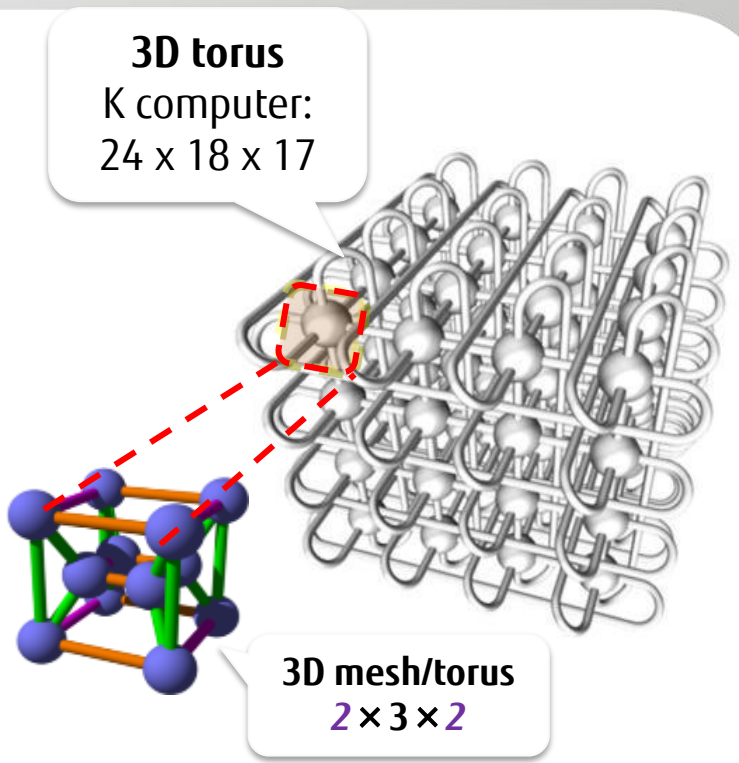

Key hardware technologies for Post-FX10

FUJITSU

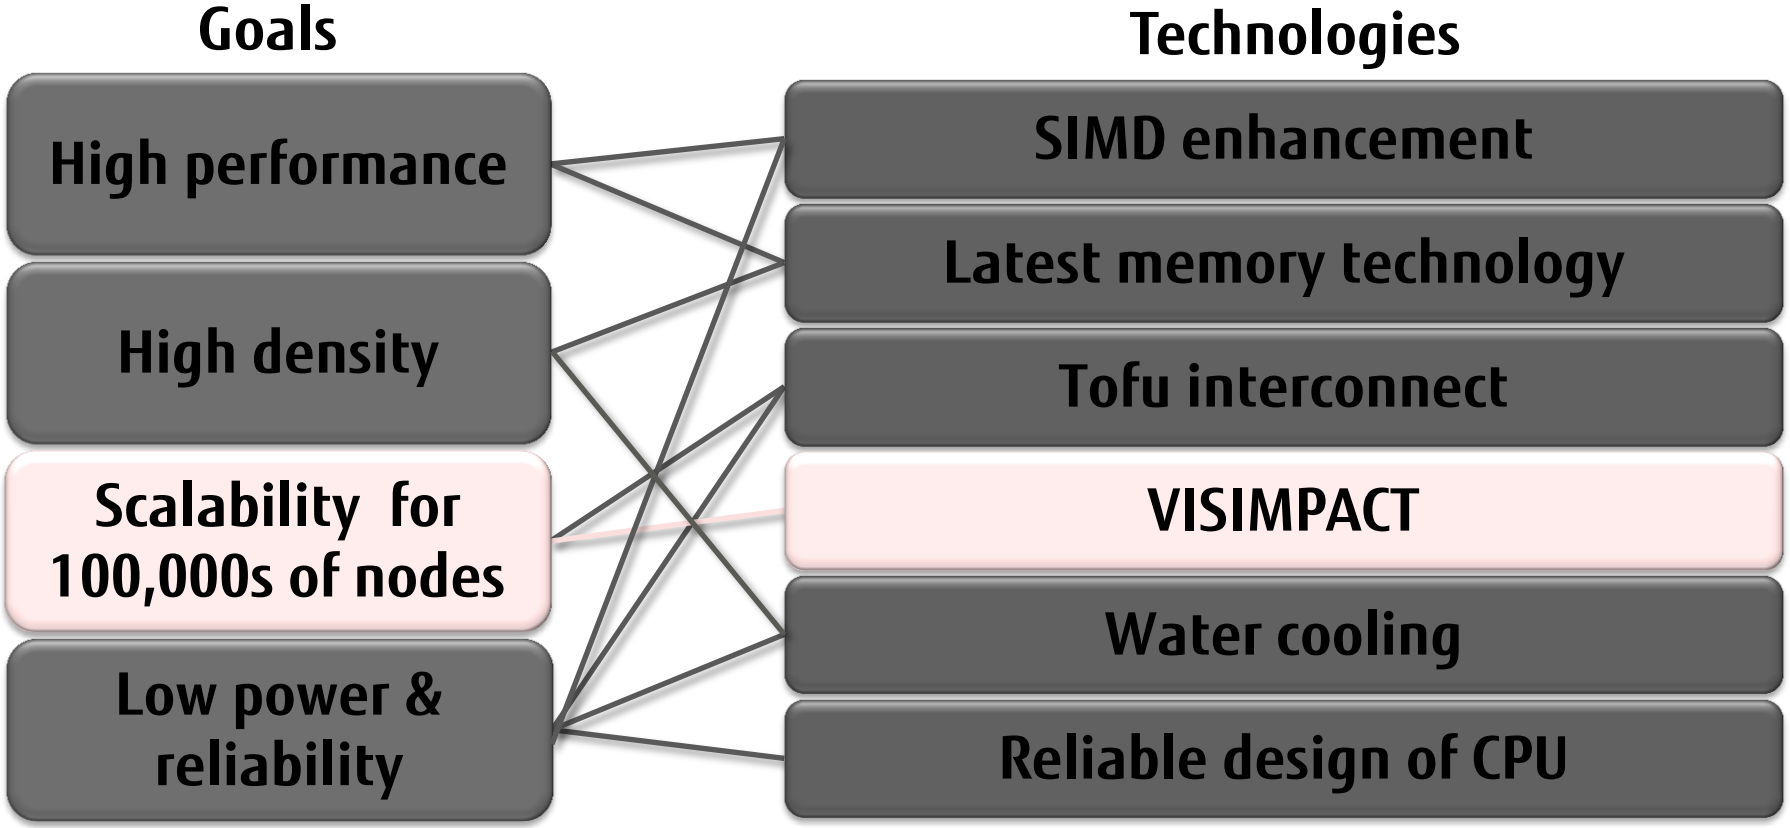

VISIMPACT:
Virtual Single Processor by Integrated Parallel Multi-core Architecture

FUJITSU

- An efficient hybrid parallel execution model and infrastructure
  - Automatic thread parallelization of MPI programs using Fujitsu compilers
  - Hardware assistance of inter-core hardware barrier and shared L2 cache
- Scalability improvement by reducing # of processes
- Increasing available memory per process

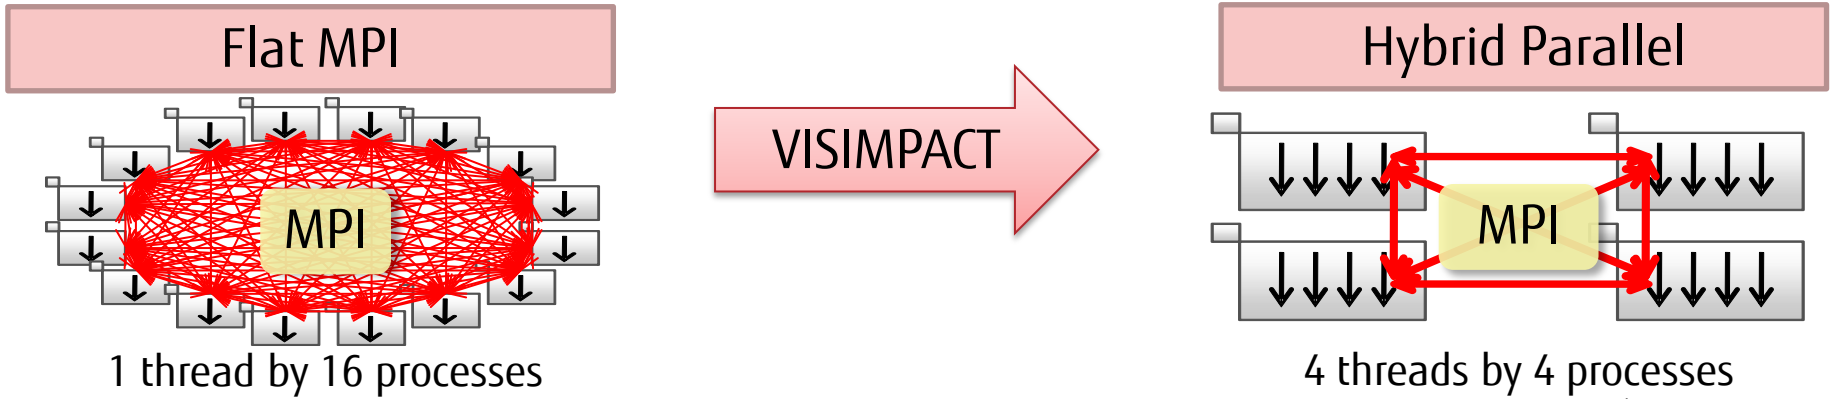

Key hardware technologies for Post-FX10

FUJITSU

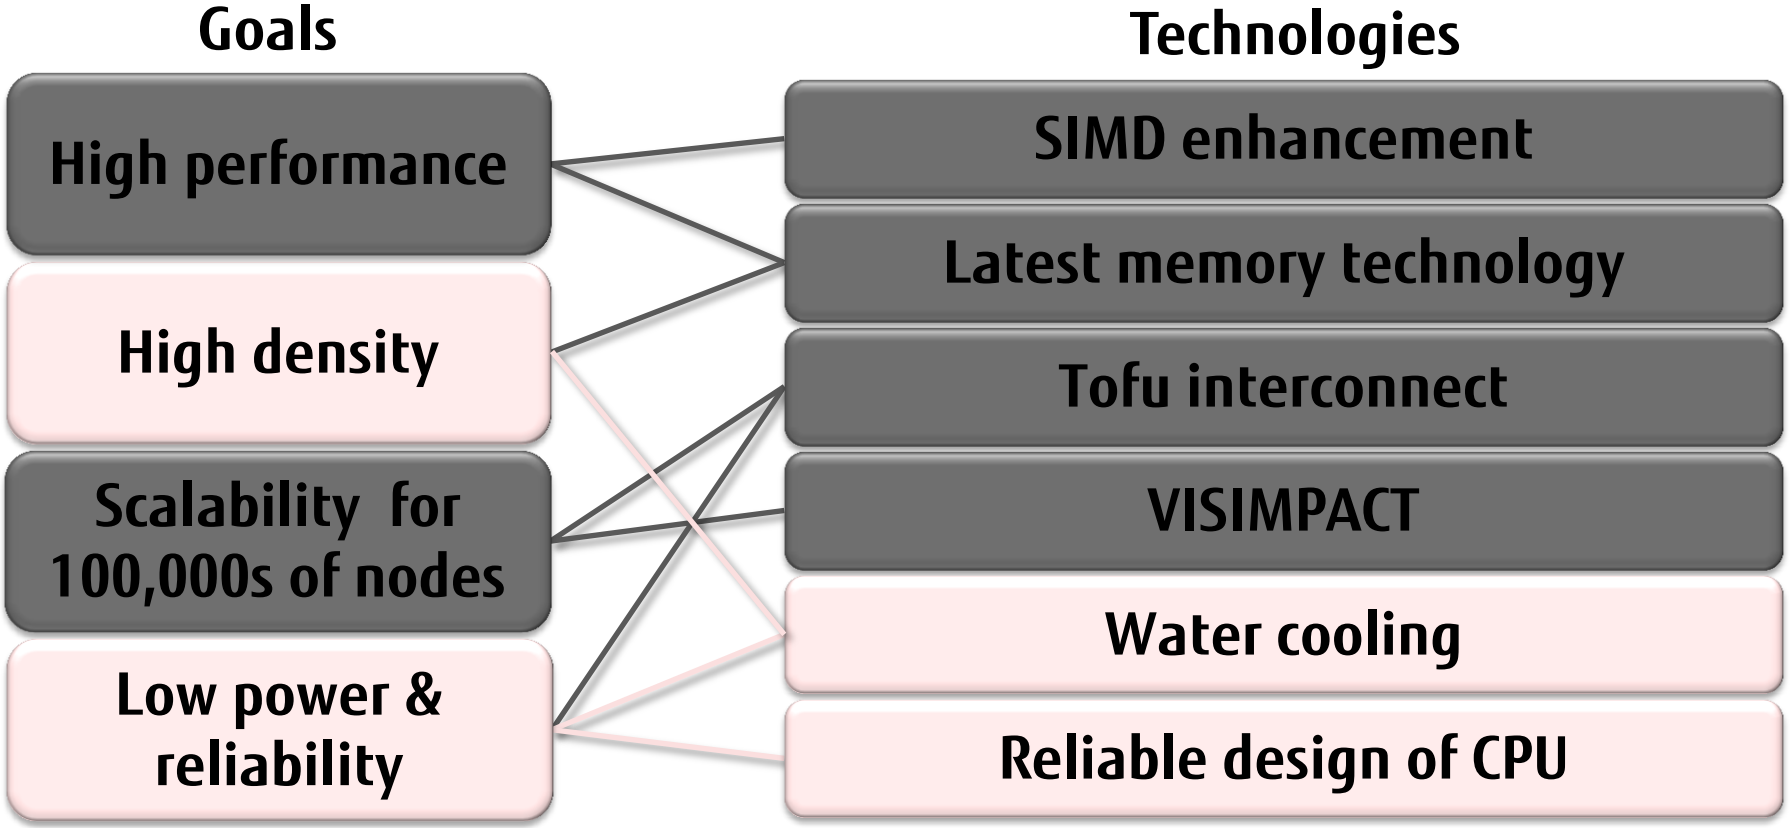

Water cooling and reliable design of CPU

FUJITSU

- Water cooling
  - All key parts are cooled by water
  - Highly reliable and low power consumption
  - High density

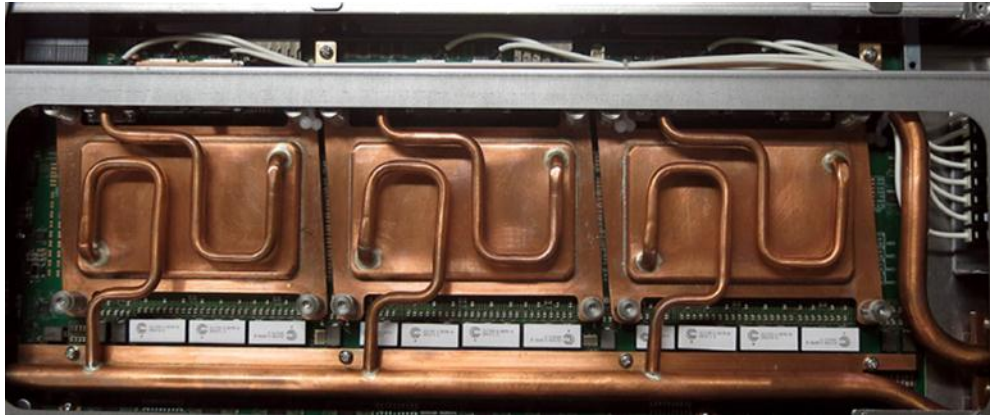

- Reliable design from mainframes
  - ECC protected L1 and L2 caches
  - Instruction retry & error recovery

SPARC64 VIIIfx

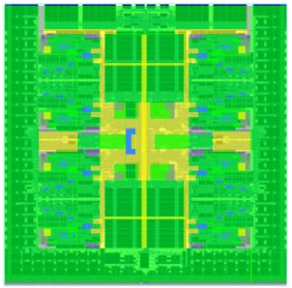

Error detection by hardware with automatic recovery

Error detection by hardware

No effect on system operation

FUJITSU

Post-FX10 prototype

- 2U chassis
- CPU memory board

Post-FX10 prototype

- 2U rack mountable chassis
  - High performance density
    - Four CPU memory boards installed
    - Total 12 nodes, over 10TFlops
  - Water cooled

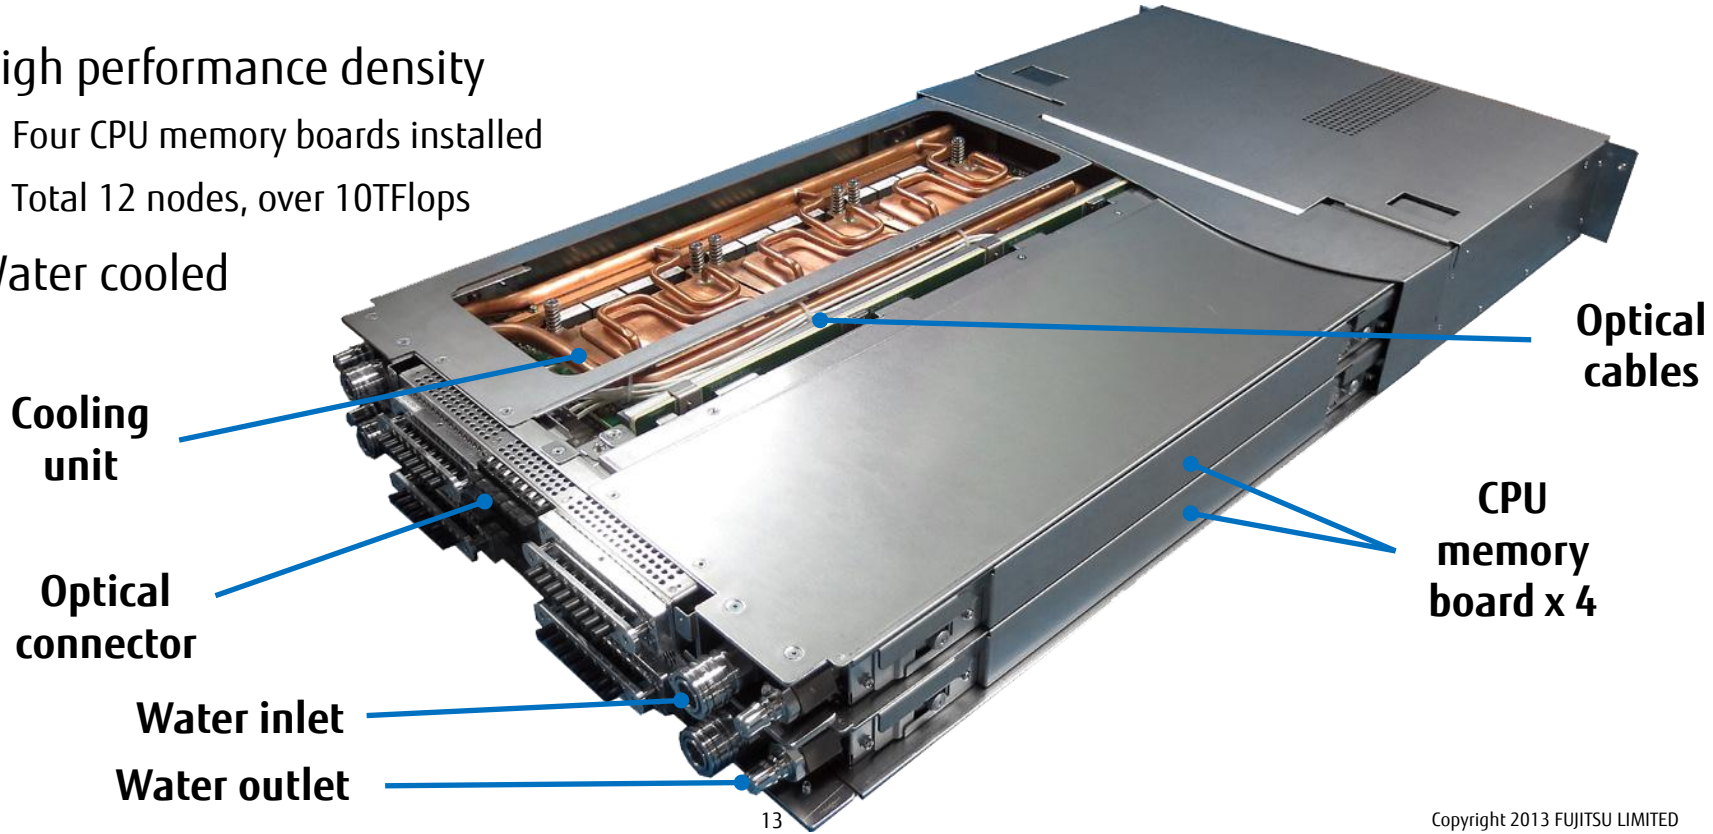

CPU memory board

- Three CPUs, nodes
  - Wide SIMD multicore
  - 1TFlops class
  - Tofu2 integrated
- HMC, Hybrid Memory Cube
  - Eight per CPU
- Optical modules

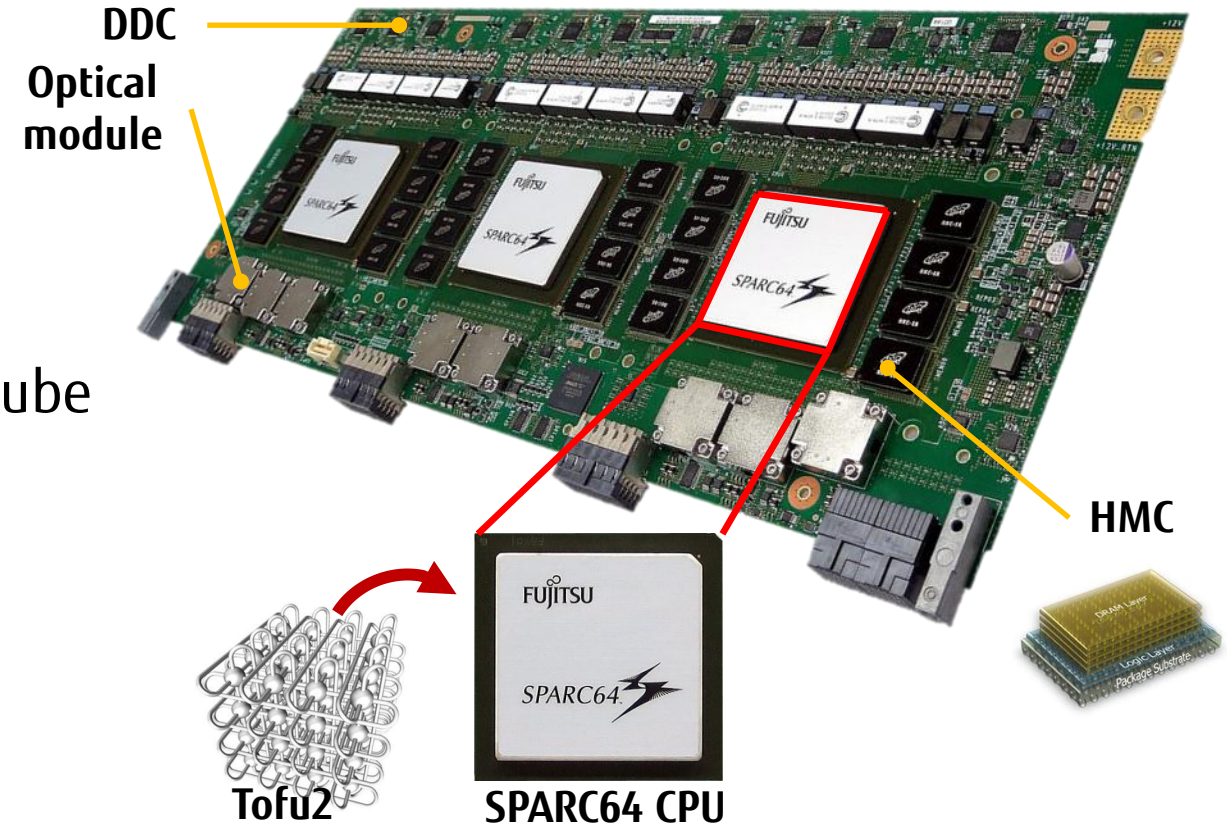

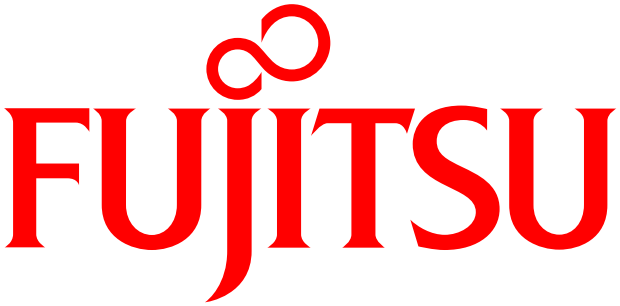

shaping tomorrow with you