FUJITSU

shaping tomorrow with you

Supercomputer “Fugaku”

Formerly known as Post-K

Supercomputer Fugaku Project

FUJITSU

RIKEN and Fujitsu are currently developing Japan's next-generation flagship supercomputer, the successor to the K computer, as the most advanced general-purpose supercomputer in the world

HPCG No.1 (2017)

AWARDS Gordon Bell Prize Finalist (2016)

GRAPH 500 No.1 (2018)

K computer

PRIMEHPC FX10

PRIMEHPC FX100

Supercomputer Fugaku

© RIKEN

- RIKEN and Fujitsu announced that manufacturing started in March 2019
- RIKEN announced on May 23, 2019 that the supercomputer is named “Fugaku”

Goals and Approaches for Fugaku

FUJITSU

Goals

High application performance

Good usability and wide range of uses

Keeping application compatibility

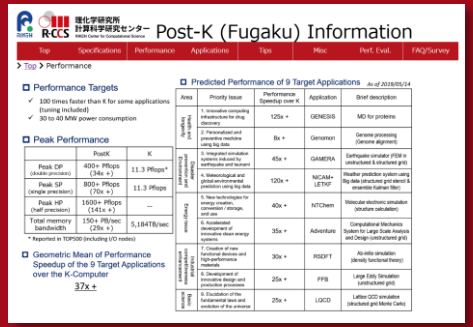
理化学研究所 計算科学研究センター Post-K (Fugaku) Information

Top > Performance

Performance Targets

- 100 times faster than K for some applications (tuning included)
- 30 to 40 MW power consumption

Peak Performance

	PostK	K
Peak DP (double precision)	400+ Pflops (34x +)	11.3 Pflops*
Peak SP (single precision)	800+ Pflops (70x +)	11.3 Pflops
Peak HP (half precision)	1600+ Pflops (141x +)	–
Total memory bandwidth	150+ PB/sec (29x +)	5,184TB/sec

* Reported in TOP500 (including I/O nodes)

Geometric Mean of Performance Speedup of the 9 Target Applications over the K-Computer

37x +

Predicted Performance of 9 Target Applications As of 2018/05/24

Area	Priority Issue	Performance Speedup over K	Application	Brief description
Health and longevity	1. Innovative computing infrastructure for drug discovery	125x +	GENESIS	MD for proteins
	2. Personalized and preventive medicine using big data	8x +	Genomon	Genome processing (Genome alignment)
Disaster prevention and global climate	3. Integrated simulation systems for hazard and disaster induced by earthquake and tsunami	45x +	GAMERA	Earthquake simulator (FEM in unstructured & structured grid)
	4. Meteorological and global environmental predictions using big data	120x +	NICAM+LETKF	Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)
Energy issue	5. New technologies for energy creation, conversion / storage, and use	40x +	NTChem	Molecular electronic simulation (structure calculation)
	6. Accelerated development of innovative clean energy systems	35x +	Adventure	Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)
Industrial competitiveness enhancement	7. Creation of new functional devices and high-performance materials	30x +	RSDFT	Ab-initio simulation (density functional theory)
	8. Development of innovative design and production processes	25x +	FFB	Large Eddy Simulation (unstructured grid)
Basic science	9. Elucidation of the fundamental laws and evolution of the universe	25x +	LQCD	Lattice QCD simulation (structured grid Monte Carlo)

RIKEN announced predicted performance:

- More than 100x+ faster than K computer for GENESIS and NICAM+LETKF
- Geometric mean of speedup over K computer in 9 Priority Issues is greater than 37x+

https://postk-web.r-ccs.riken.jp/perf.html

Goals and Approaches for Fugaku

FUJITSU

Goals

High application performance

Good usability and wide range of uses

Keeping application compatibility

Approaches

Develop

1. High-performance Arm CPU A64FX in HPC and AI areas
2. Cutting-edge hardware design
3. System software stack

Achieve

- High performance in real applications
- High efficiency in key features for AI applications

1. High-Performance Arm CPU A64FX in HPC and AI Areas

FUJITSU

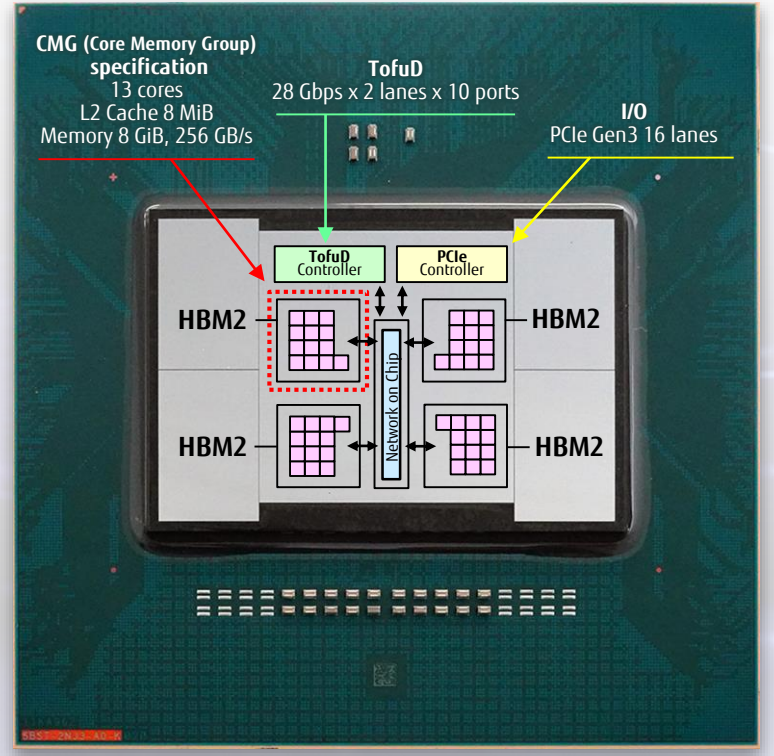

Architecture features

ISA	Armv8.2-A (AArch64 only) SVE (Scalable Vector Extension) arm
SIMD width	512-bit
Precision	FP64/32/16, INT64/32/16/8
Cores	48 computing cores + 4 assistant cores (4 CMGs)
Memory	HBM2: Peak B/W 1,024 GB/s
Interconnect	TofuD: 28 Gbps x 2 lanes x 10 ports

Peak performance (Chip level)

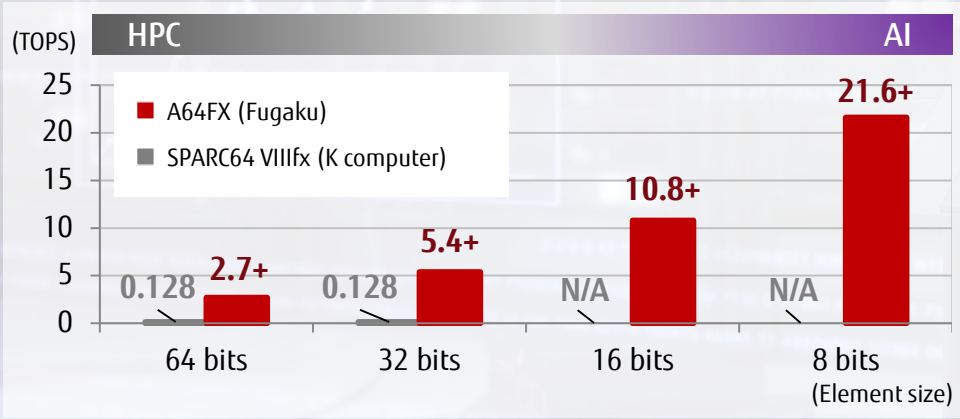

1. High-Performance Arm CPU A64FX in HPC and AI Areas

FUJITSU

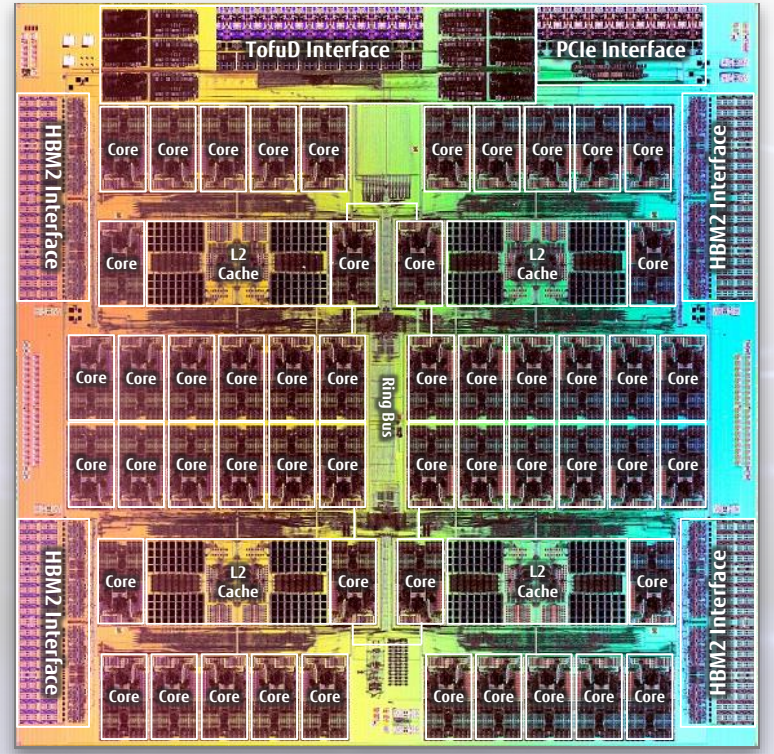

Architecture features

ISA	Armv8.2-A (AArch64 only) SVE (Scalable Vector Extension) arm
SIMD width	512-bit
Precision	FP64/32/16, INT64/32/16/8
Cores	48 computing cores + 4 assistant cores (4 CMGs)
Memory	HBM2: Peak B/W 1,024 GB/s
Interconnect	TofuD: 28 Gbps x 2 lanes x 10 ports

Peak performance (Chip level)

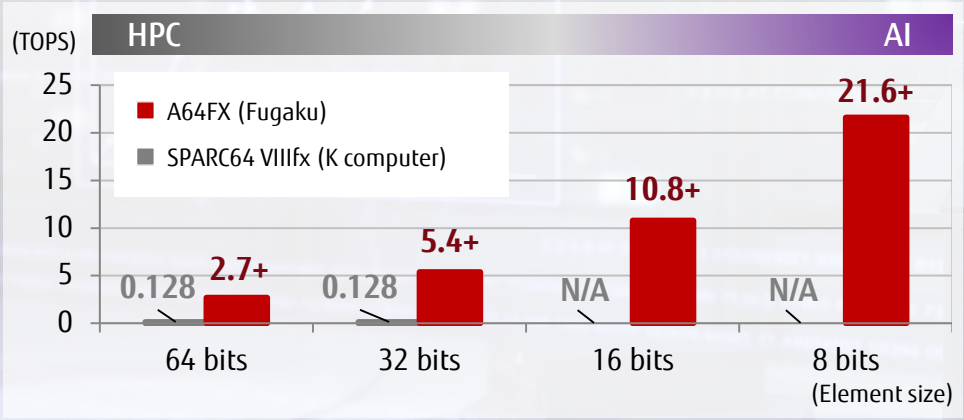

2. Cutting-edge Hardware Design

FUJITSU

1PFlops by Fugaku and K computer

	Fugaku	K computer
Configuration	1x rack including SSDs	80x compute racks & 20x disk racks
Nodes	384	8,160
Footprint	1.1 m^2 (0.8 m x 1.4 m)	128 m^2 (4 m x 32 m)

Scalable design

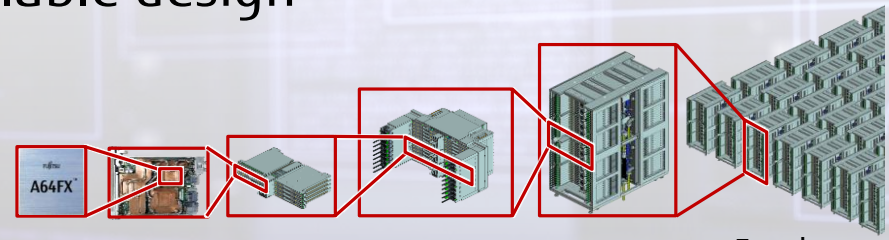

	CPU	CMU	BoB	Shelf	Rack	Fugaku
Nodes	1	2	16	48	384	150k+
Performance [Flops]	2.7 T+	5.4 T+	43 T+	129 T+	1 P+	400 P+

CMU (CPU Memory Unit)

- 100% direct water cooling
- 3x QSFP for AOC (Active Optical Cables)
- Single-sided blind mate connectors for electrical signals and water

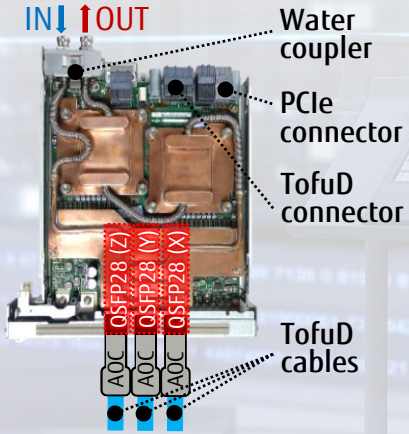

3. System Software Stack

FUJITSU

- Fujitsu developing system software in collaboration with RIKEN
  - Fujitsu Technical Computing Suite implementing development and execution environments with great usability on large-scale system

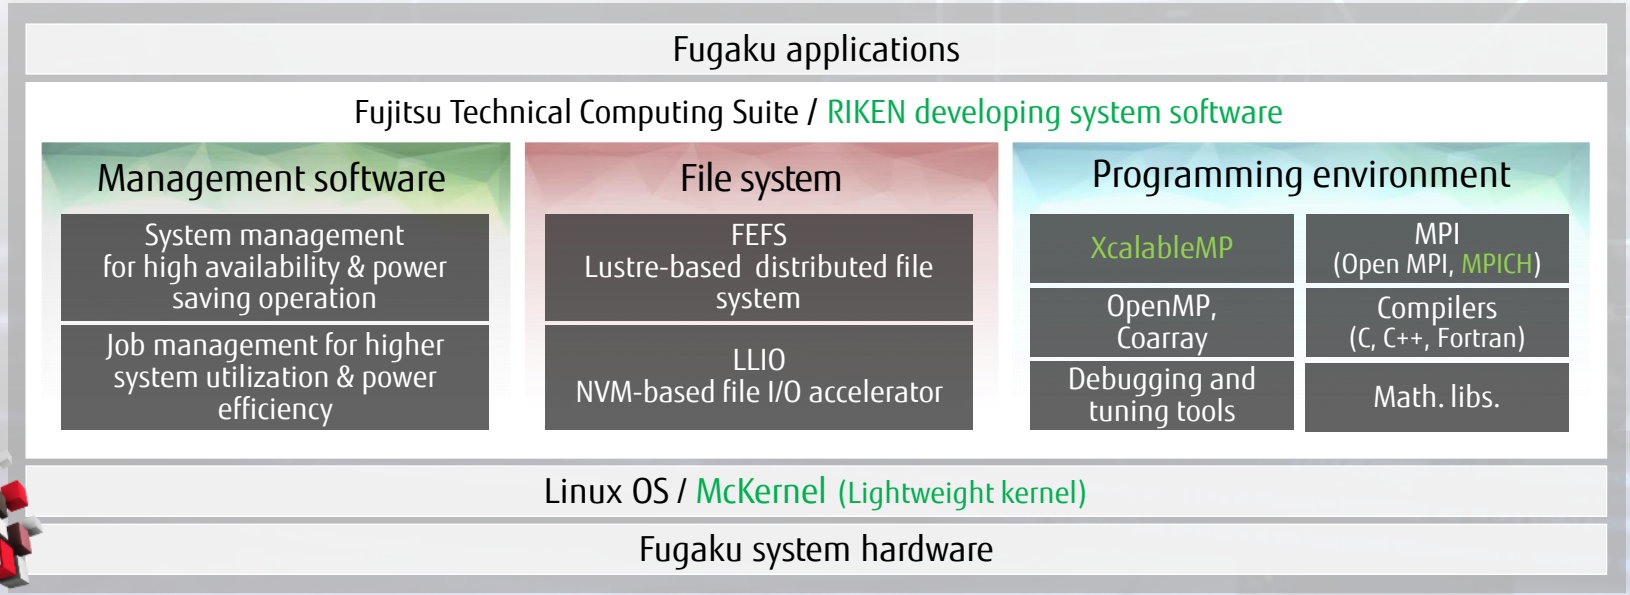

3. System Software Stack

FUJITSU

- Fujitsu developing system software in collaboration with RIKEN
  - Fujitsu Technical Computing Suite implementing development and execution environments with great usability on large-scale system

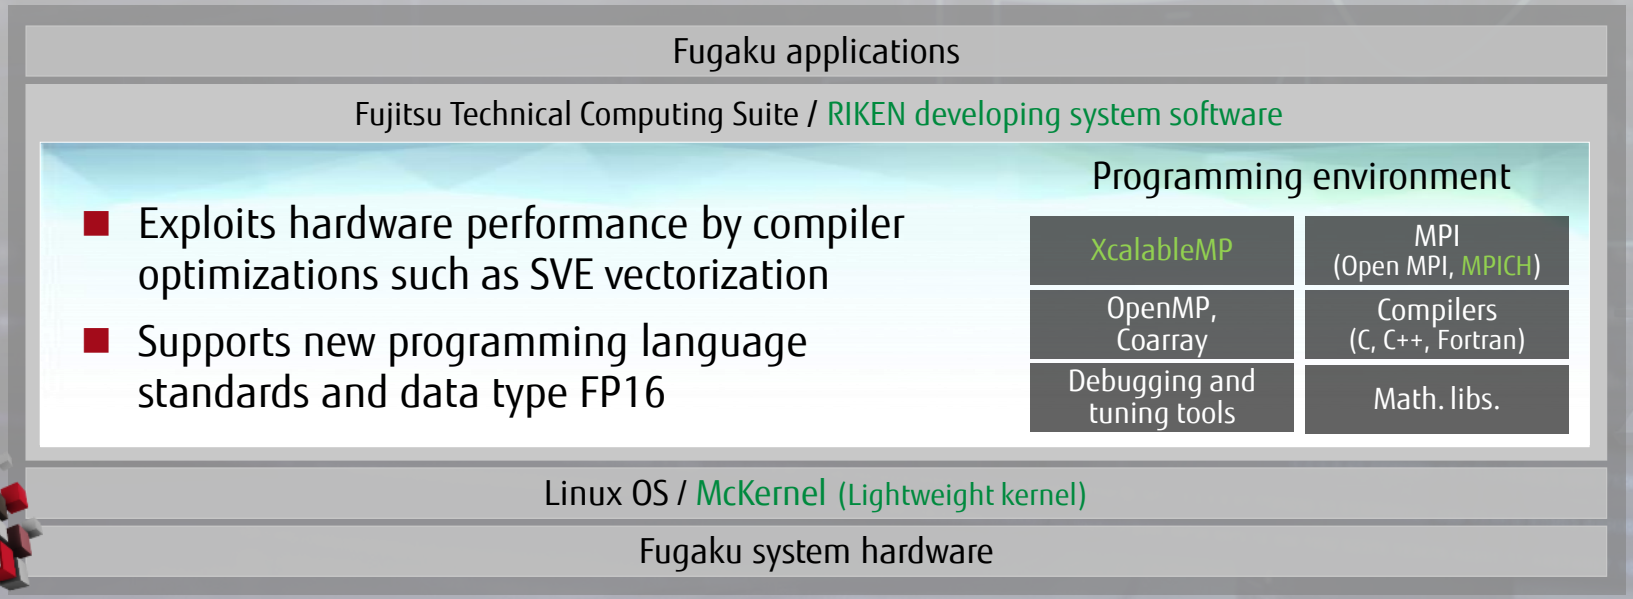

Fugaku Under development w/ RIKEN

Achieve

High Performance in Real Application

FUJITSU

- WRF: Weather Research and Forecasting model (v3.8.1)
  - Vectorizing loops including IF statements is key optimization
- Himeno Benchmark (Fortran90, size: XL)
  - Stencil calculation to solve Poisson's equation by Jacobi method

High memory B/W and long SIMD length work effectively

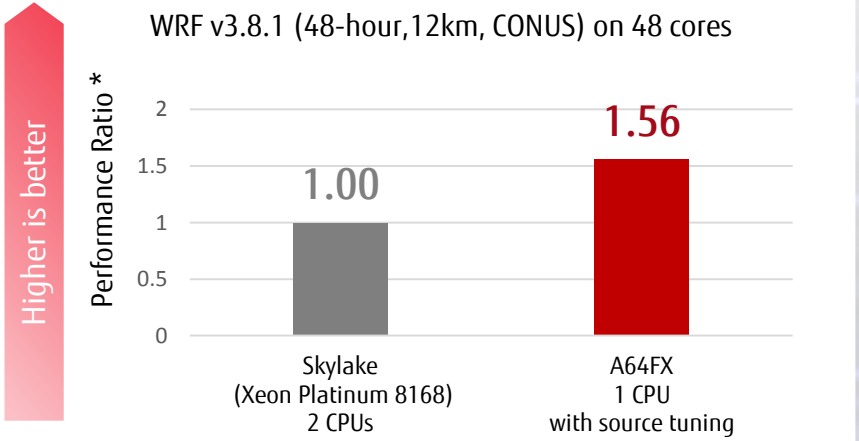

* Normalized by the average elapsed time for timestep of Skylake

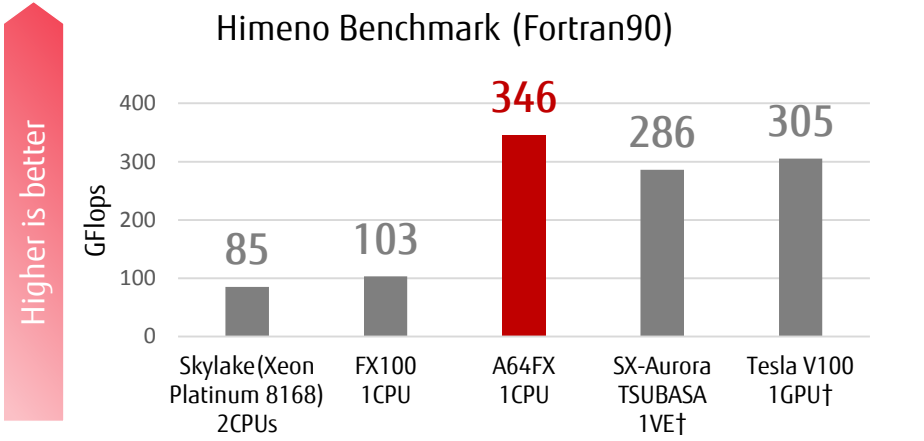

†Performance evaluation of a vector supercomputer SX-aurora TSUBASA https://dl.acm.org/citation.cfm?id=3291728

Achieve

High Efficiency in Key Features for AI Applications

FUJITSU

High FP16 & INT8 peak performance and high memory peak B/W

FP16 performance: **10.8+ TOPS, > 90%@HGEMM**

INT8 performance : **21.6+ TOPS in partial dot product**

Memory B/W : **1,024 GB/s, > 80%@STREAM Triad**

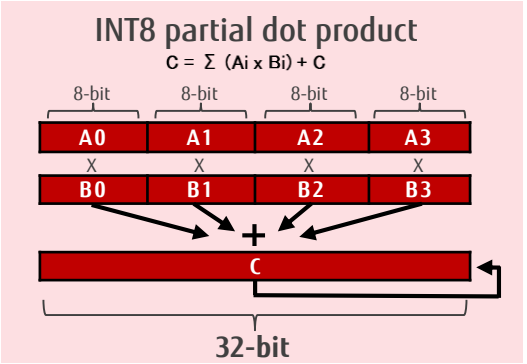

Functions contributing to key features in AI fields

A64FX CPU	• 2x 512-bit wide SIMD pipelines per core for FP16 and INT8 • High memory B/W and calculation throughput
Compilers & libraries	• Vectorization and software pipelining • FP16 as data type of programming language (e.g., real (kind=2) in Fortran) • Mathematical Library for HGEMM

Future Plans

FUJITSU

Supercomputer Fugaku

- Operations starting around CY2021

Today

CY2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Fugaku				Basic Design	Design Implementation		Manufacturing, Installation and Tuning		Operations	
Development				Operations			Operations ending Aug, 2019			

© RIKEN

K computer

Fujitsu HPC Products

- Fujitsu will begin global sales of supercomputers based on the Supercomputer Fugaku technology in the 2nd half of FY2019

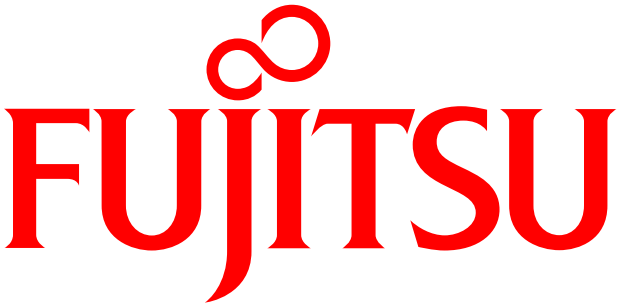

shaping tomorrow with you