

SPARC64™ X+: **Fujitsu's Next Generation Processor** **for UNIX servers**



August 27, 2013

Toshio Yoshida

Processor Development Division
Enterprise Server Business Unit
Fujitsu Limited

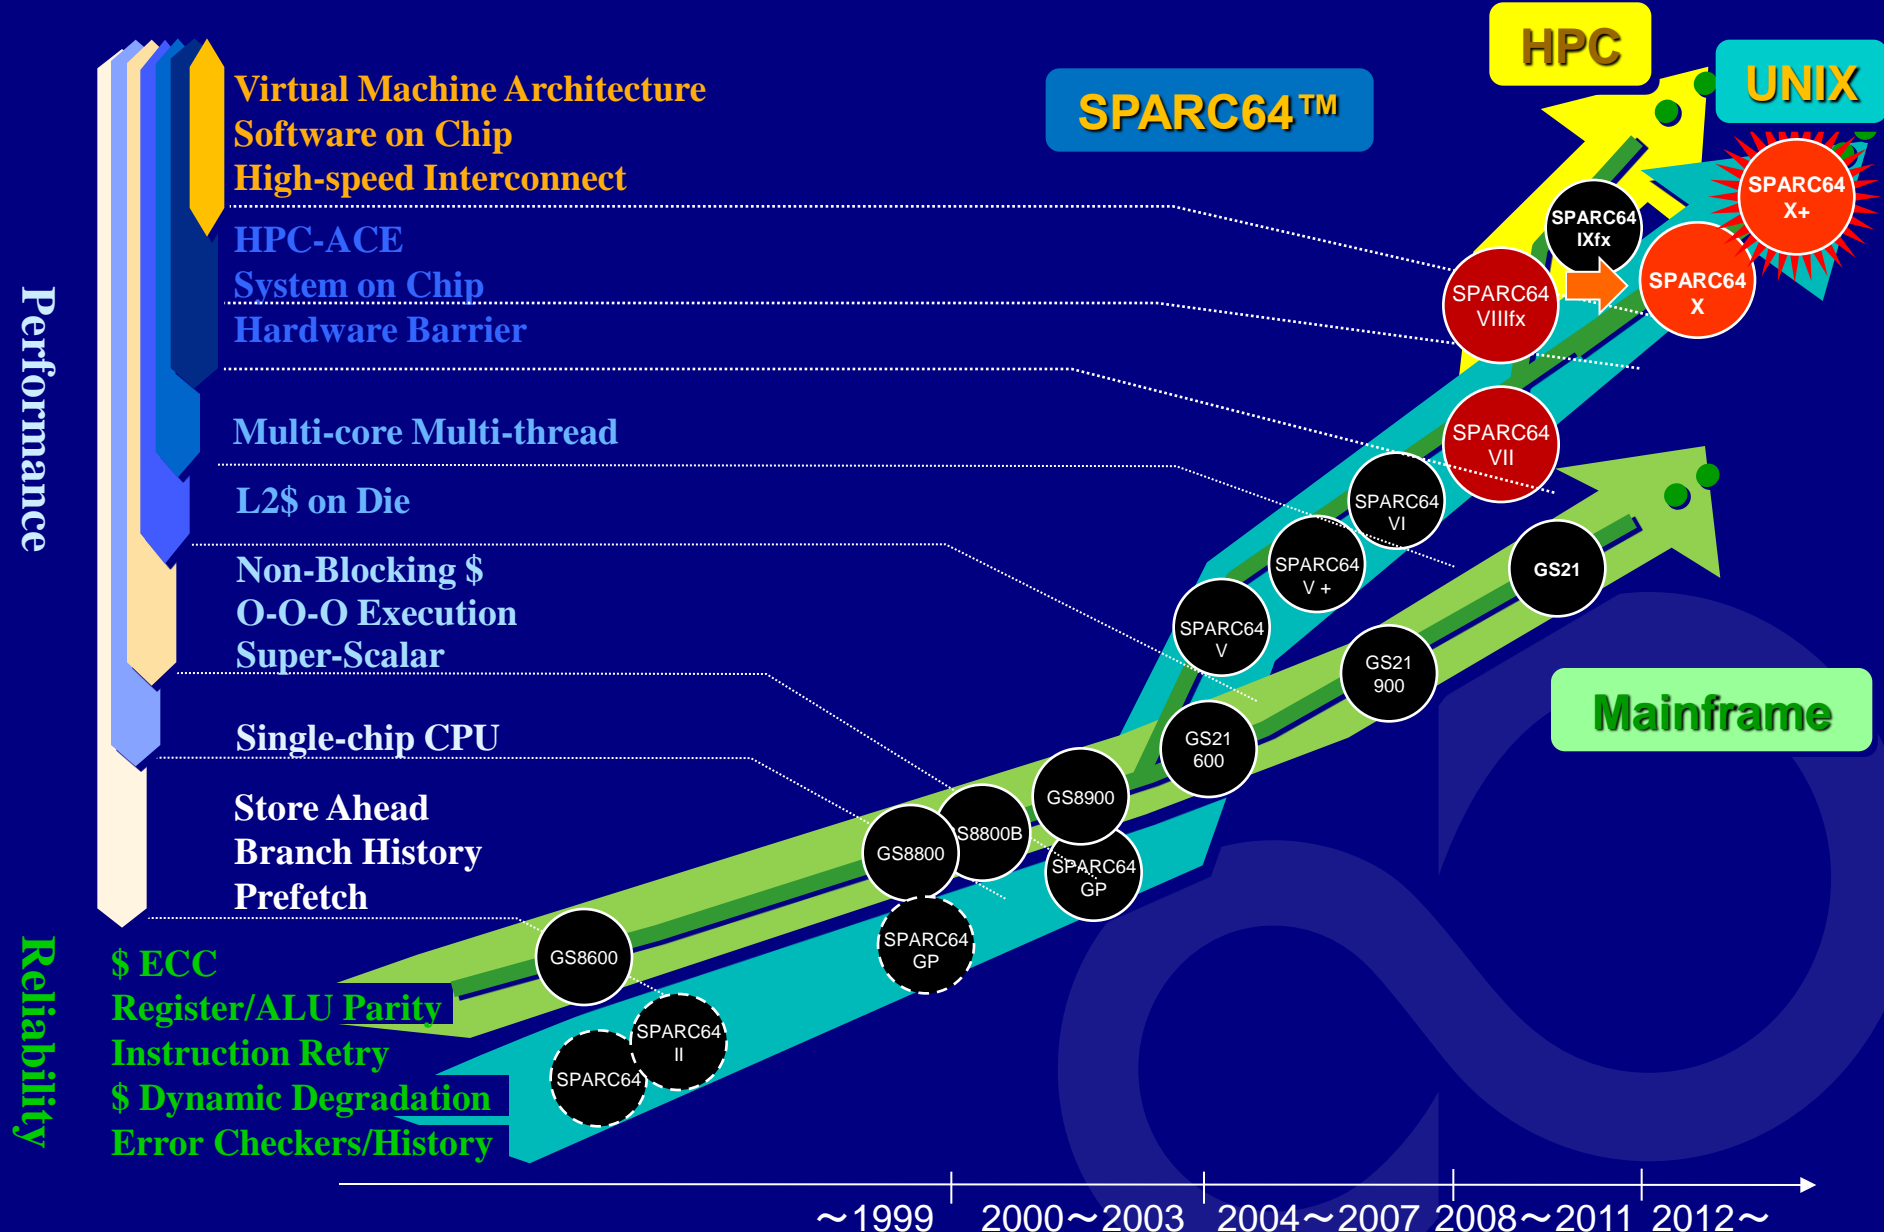
Agenda

- ◆ Fujitsu Processor Development

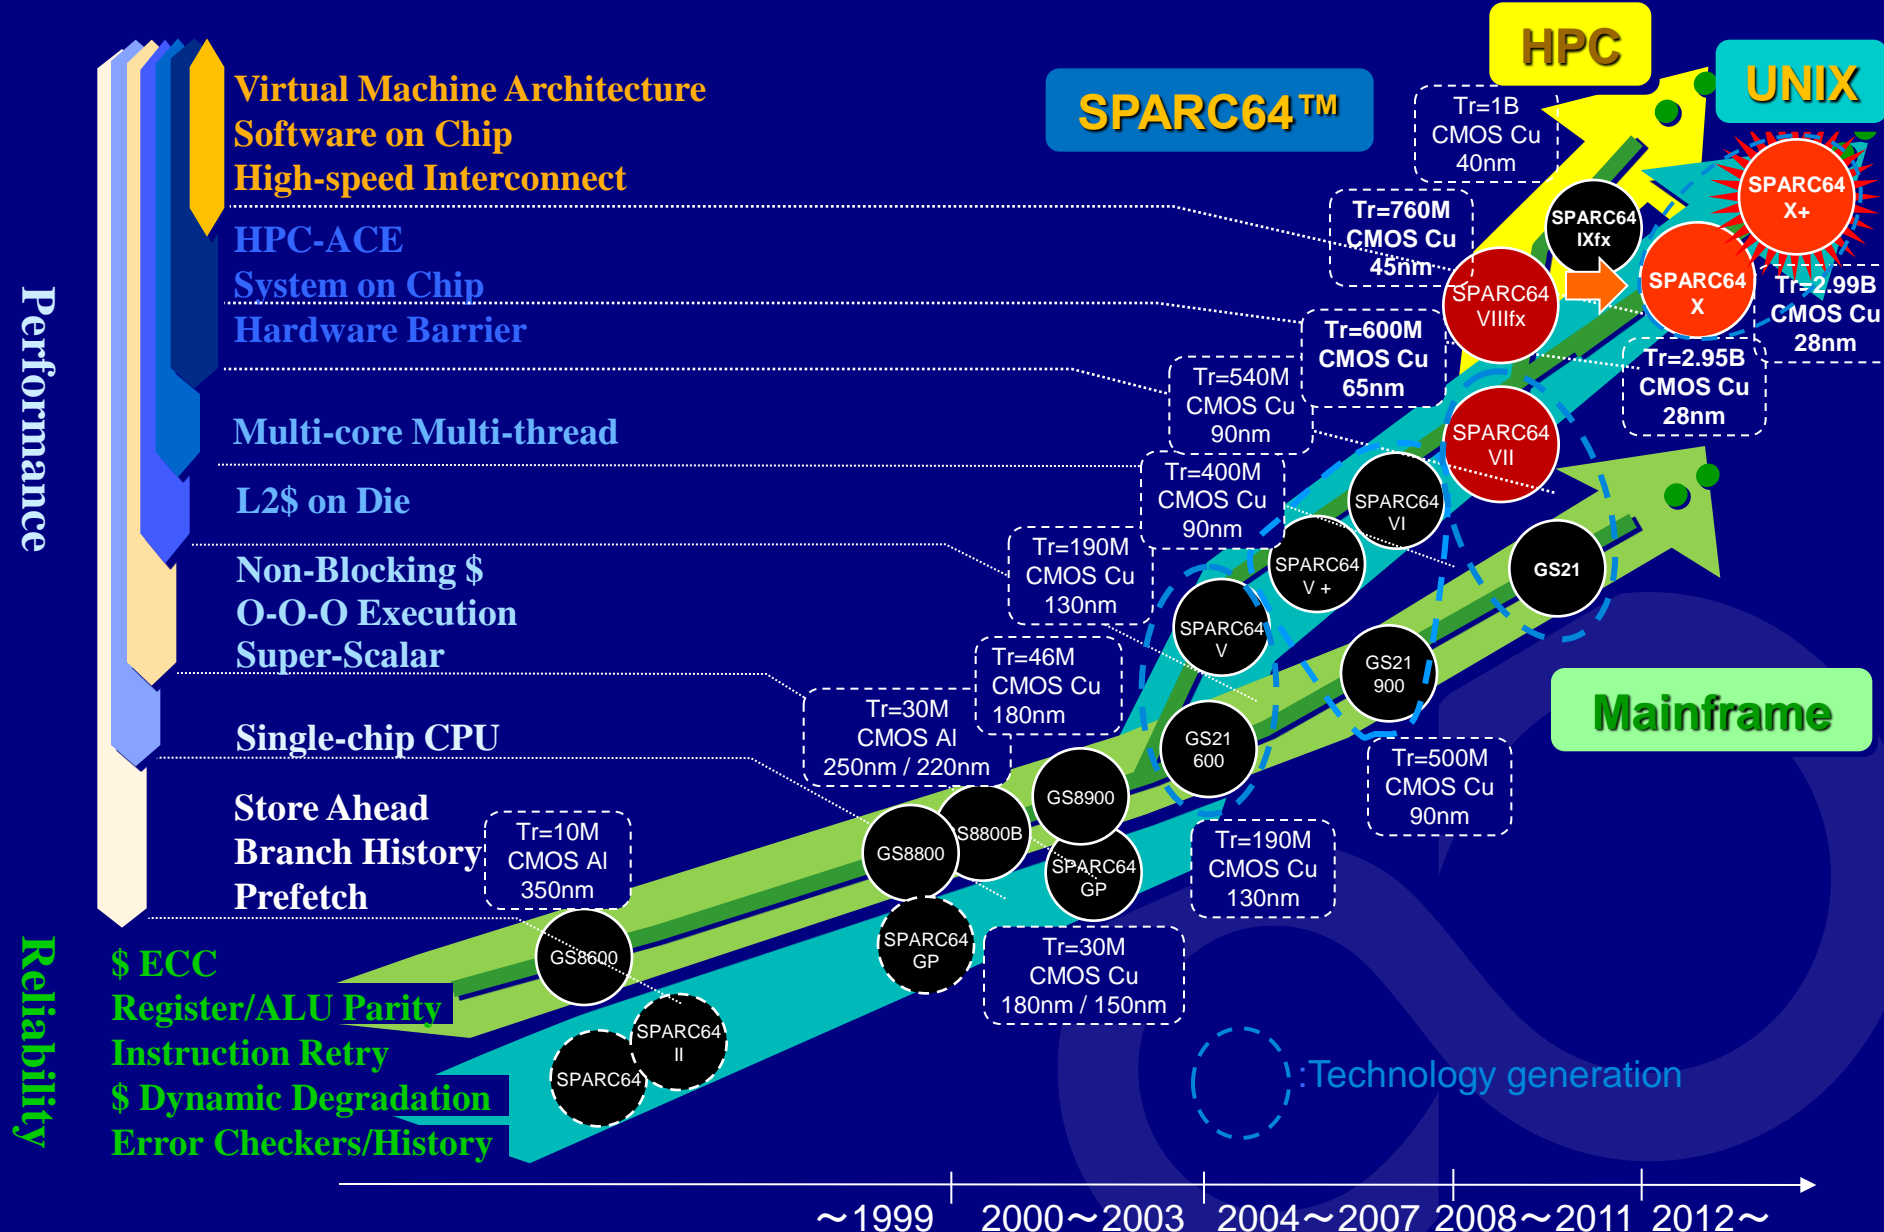
- ◆ SPARC64™ X+
 - Design Concept and Processor Overview
 - Software on Chip (SWoC)
 - Micro-Architecture
 - System Architecture
 - RAS
 - Power Management

- ◆ Summary

Fujitsu Processor Development



Fujitsu Processor Development



SPARC64™ X+

Agenda

- ◆ Fujitsu Processor Development

- ◆ SPARC64™ X+
 - Design Concept and Processor Overview
 - Software on Chip (SWoC)
 - Micro-Architecture
 - System Architecture
 - RAS
 - Power Management

- ◆ Summary



Design of SPARC64™ X / X+

◆ Combine Fujitsu HPC and UNIX processor features

✓ Single-Thread Performance

- Higher clock speed
- Micro-architectural enhancements
- Directly connected DIMMs

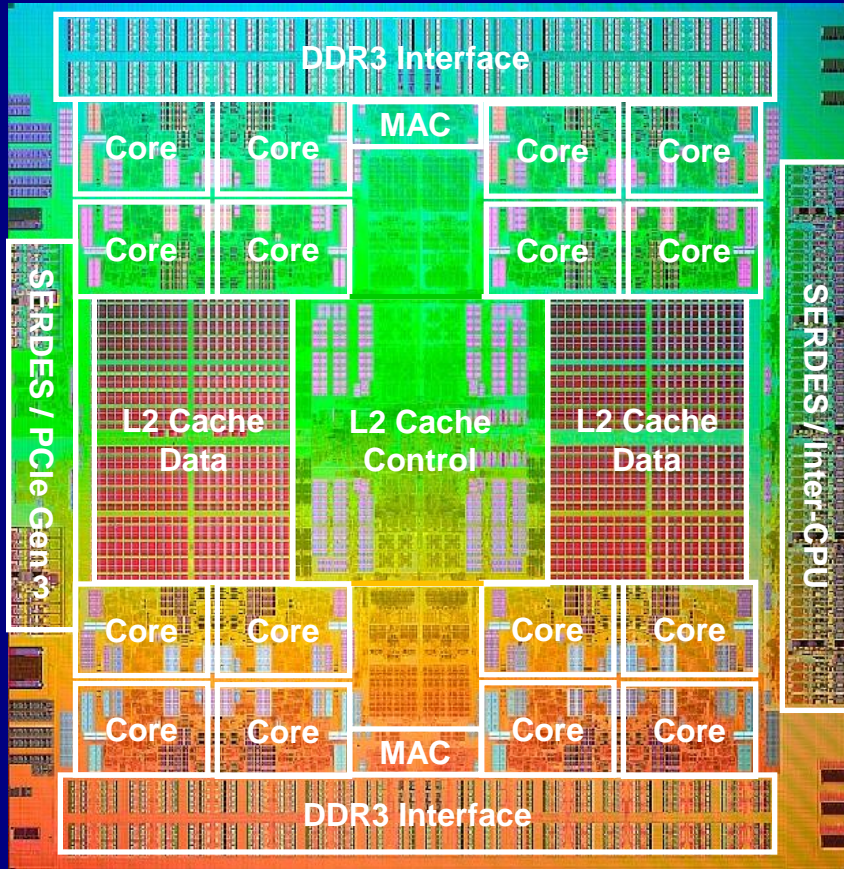
✓ High Throughput for massive data processing

- SIMD parallelism and more registers
- Multi-core and multi-thread
- High bandwidth interconnect and memory links
- Scalability up to 64 sockets (2048 threads)

✓ Software on Chip (SWoC) for specific applications

- Cipher, Decimal, Database

SPARC64™ X+ Chip Overview



● Architecture Features

- 16 cores x 2 SMT threads
- Shared 24 MB L2\$
- Memory and I/O Controllers
- HPC-ACE
- SWoC (Software on Chip)

● 28nm CMOS

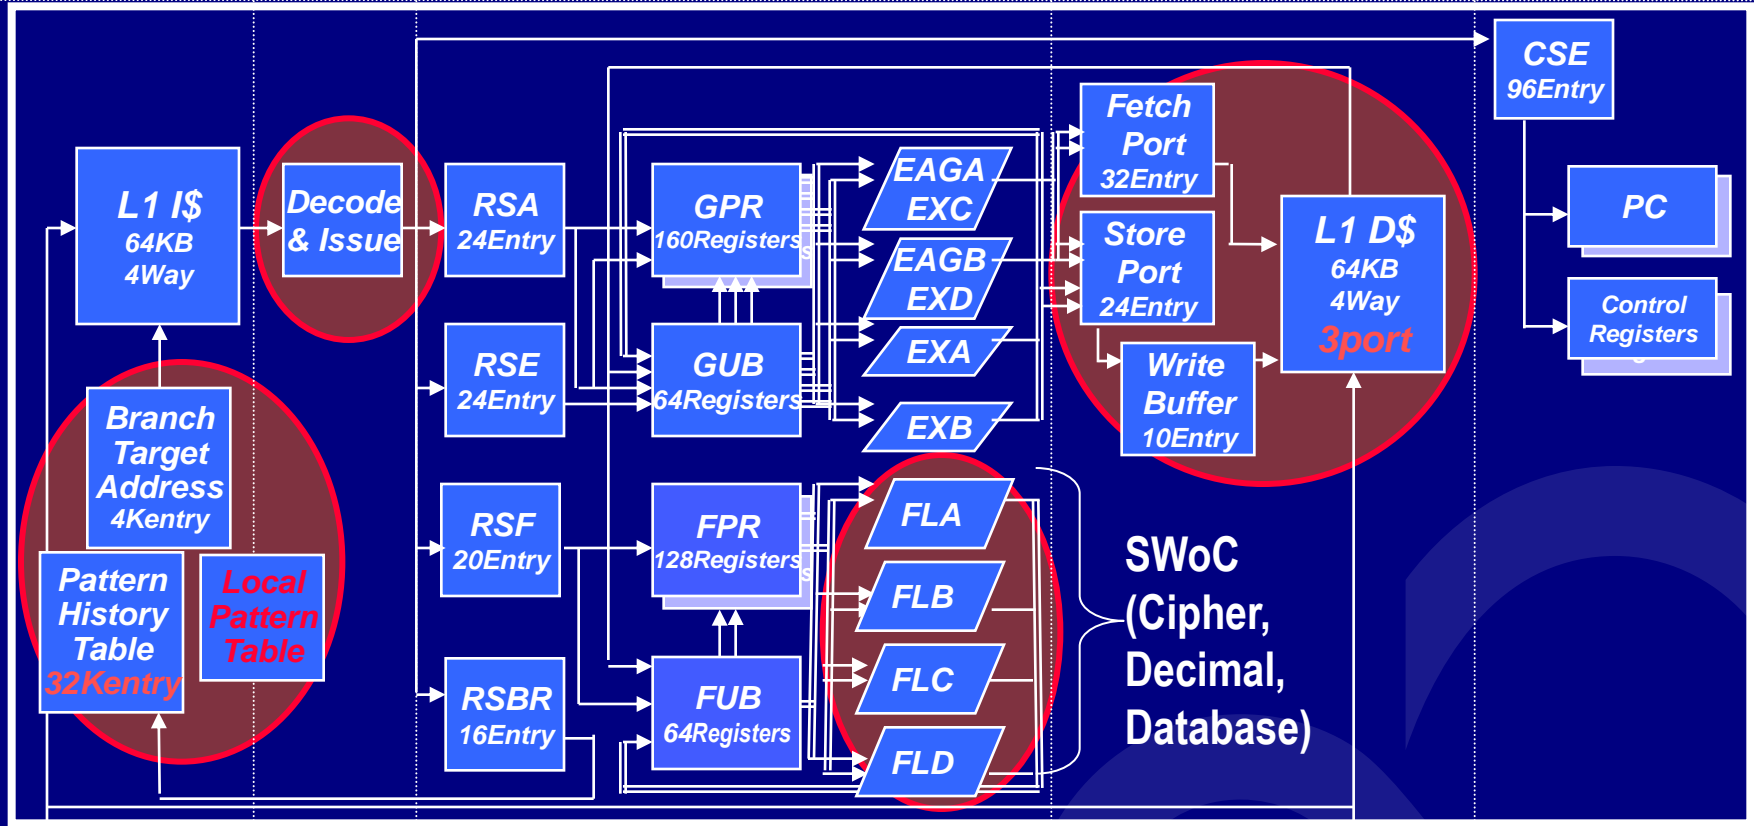
- 24.0mm x 25.0mm
- 2,990M transistors
- 1,500 signal pins
- 3.5GHz+

● Performance (peak)

- 448GFlops+
- 102GB/s memory throughput

SPARC64™ X+ Pipeline

Fetch (4 stages) Issue (4 stages) Dispatch (5 stages) Reg-Read (5 stages) Execute (5 stages) Memory (L1\$: 3 stages) Commit (2 stages)



Enhanced in SPARC64™ X+

Router
25Gbps x 3ports
14/14.5Gbps x 2ports
CPU-CPU I/F

L2\$
24MB 24Way
Memory Controller
DIMM

16-core
IO Controller
PCI-GEN3

Agenda

◆ Fujitsu Processor Development

◆ SPARC64™ X+

- Design Concept and Processor Overview
- Software on Chip (SWoC)
- Micro-Architecture
- System Architecture
- RAS
- Power Management

◆ Summary



Software on Chip (SWoC)


◆ SPARC64™ X / X+ Software on Chip

- ✓ Cipher
- ✓ Decimal (IEEE754 DPD, NUMBER)
- ✓ Database processing

◆ Accelerate specific software functions in hardware

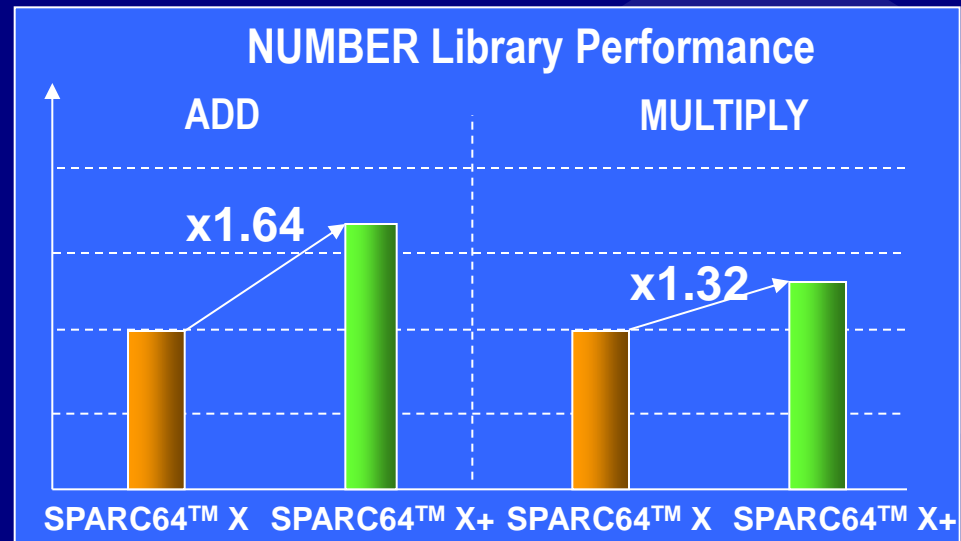
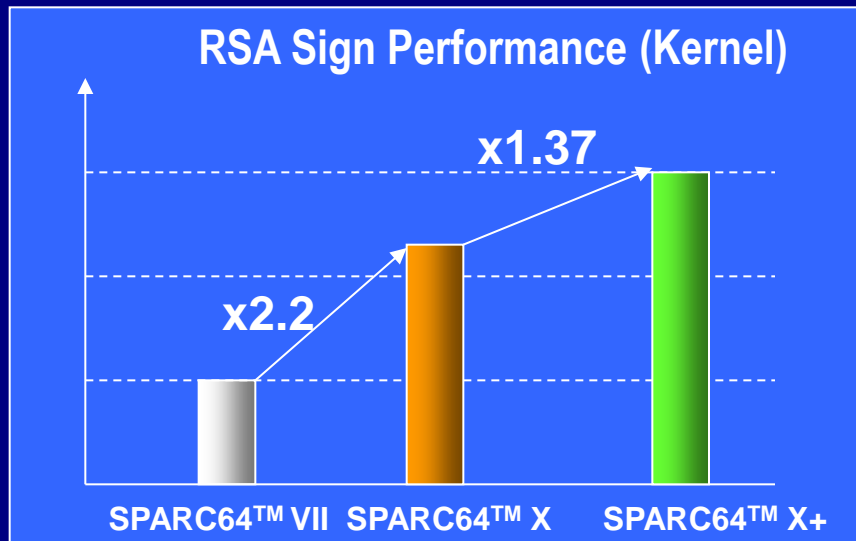
- ✓ SWoC engines implemented in floating-point unit can use 128 floating-point registers, software pipelining
- ✓ Area/number of gates < 3% of core and < 1% of chip

SPARC64™ X Core



Cipher and Decimal Performance

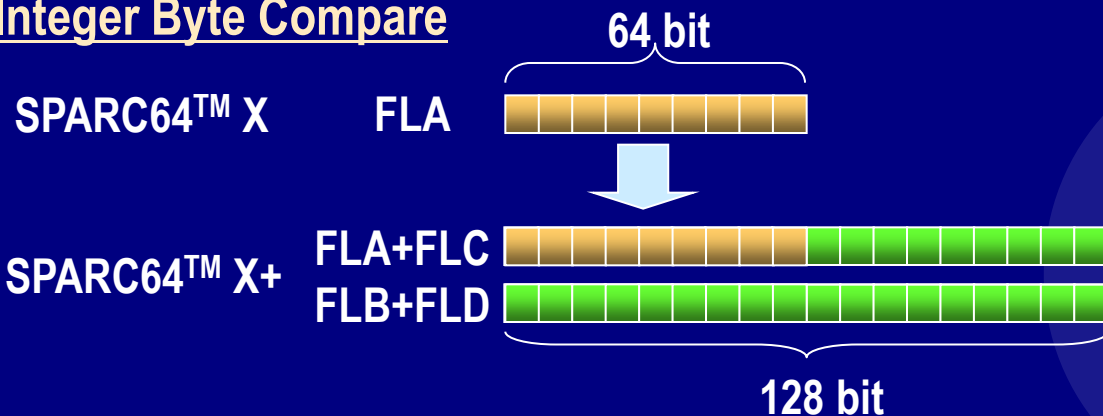
- Cipher
 - AES/DES/SHA/RSA in SPARC64™ X
 - RSA further improved in SPARC64™ X+
 - New instruction for RSA sign library
- Decimal
 - SPARC64™ X+ micro-architectural enhancements speed up several NUMBER libraries



Database Acceleration

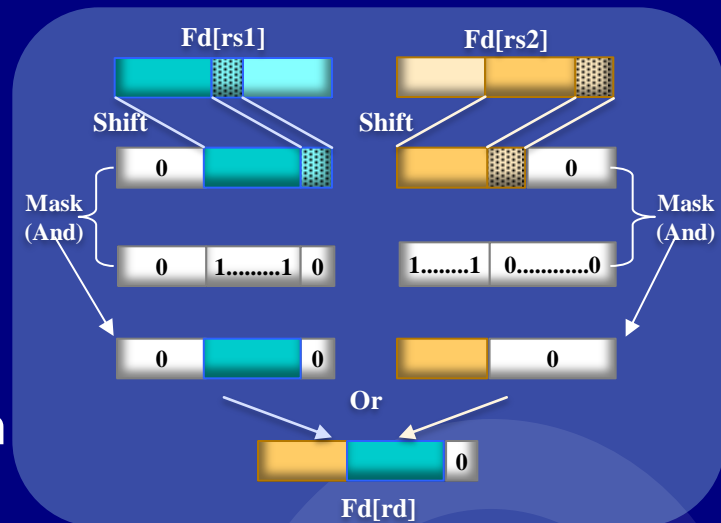
- Fine-grained data manipulation
 - Byte vector in SPARC64™ X
 - Bit vector enhanced in SPARC64™ X+
- Integer Byte Compare
 - Enhanced ISA supports SIMD operation
 - Enhanced core supports instruction in both floating-point pipelines

Integer Byte Compare

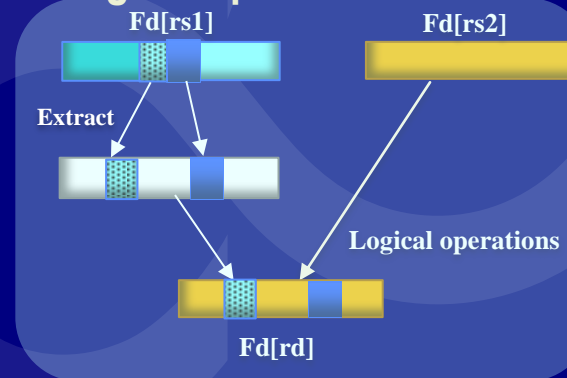


Bit Vector Operations

Shift -> Mask -> Or



Extract 2 bit fields from rs1
-> Logical operation with rs2



Agenda

◆ Fujitsu Processor Development

◆ SPARC64™ X+

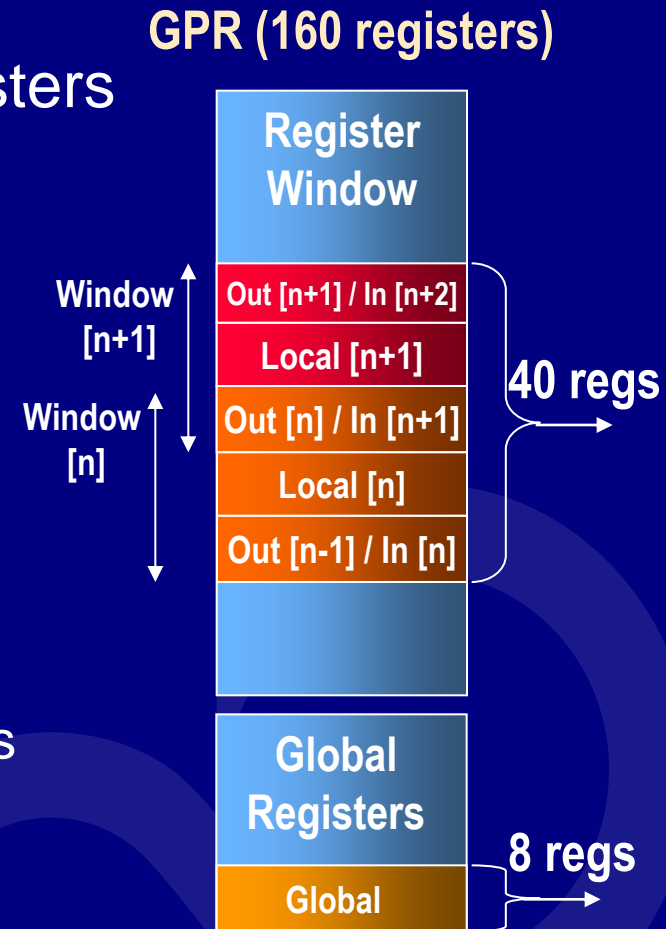
- Design Concept and Processor Overview
- Software on Chip (SWoC)
- Micro-Architecture
- System Architecture
- RAS
- Power Management

◆ Summary



Micro-Architectural Enhancements 1/2


- Register window switches
 - Out-of-order access to 48 integer registers (current & next window)
 - No penalty for all window switches between same two windows
 - SPARC64™ X handles only one window switch without penalty
- Improved Branch prediction
 - Rehashed indirect branch predictor
 - Indirect branch with variable target address
 - Local pattern branch predictor
 - More pattern history table entries



Micro-Architectural Enhancements 2/2

- L1 data cache
 - Dedicated write pipeline
 - 64 RAM banks (8 sets of 8-banked RAMs)
 - One write and two reads each cycle, except when RAM bank conflict occurs
 - Faster atomic memory operations
 - Increased hardware prefetch throughput

L1-D Cache Schematic



L1-D Cache Throughput



Agenda

◆ Fujitsu Processor Development

◆ SPARC64™ X+

- Design Concept and Processor Overview
- Software on Chip (SWoC)
- Micro-Architecture
- System Architecture
- RAS
- Power Management

◆ Summary



SPARC64™ X / X+ System Architecture

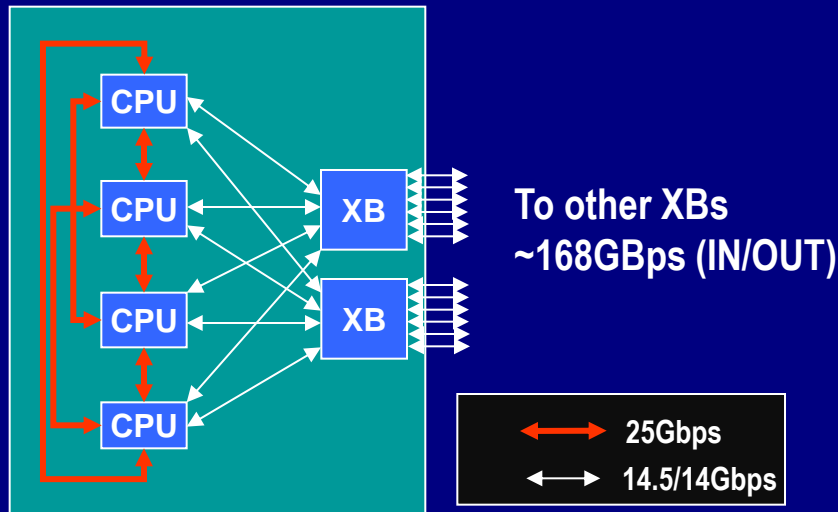
◆ Scales from 1 to 64 CPU sockets (2048 threads)

- Directory-based cache coherency
- High-speed interconnect, up to 25Gbps per lane in SPARC64™ X+ (Up to 14.5Gbps in SPARC64™ X)

◆ System Configuration

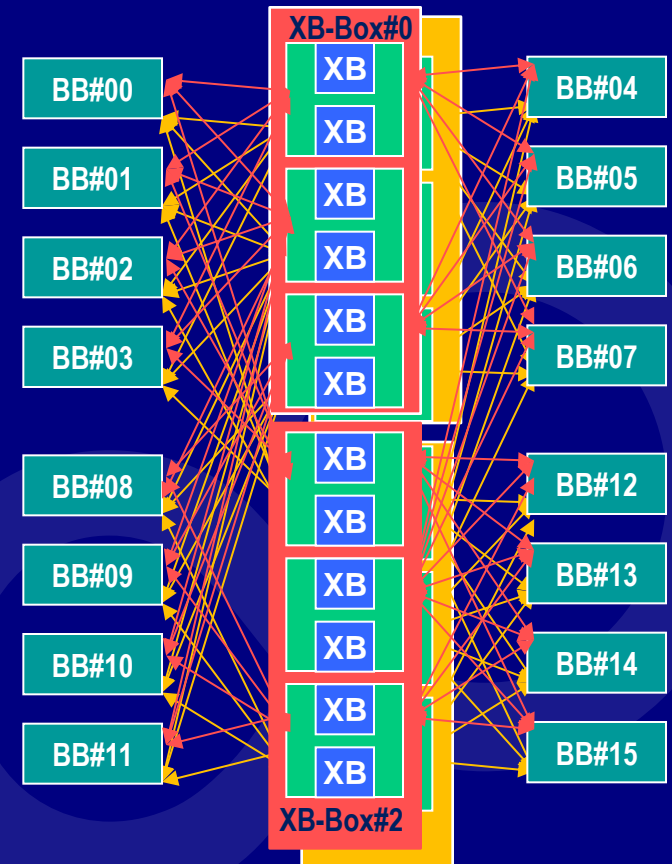
- Building Block (BB) is 4 CPUs and 2 XBs
- Up to 4 BBs can be connected by XBs
- 16BBs can be connected via XB-Boxes

Building Block (4 CPU Sockets)



16 BBs (64 CPU Sockets)

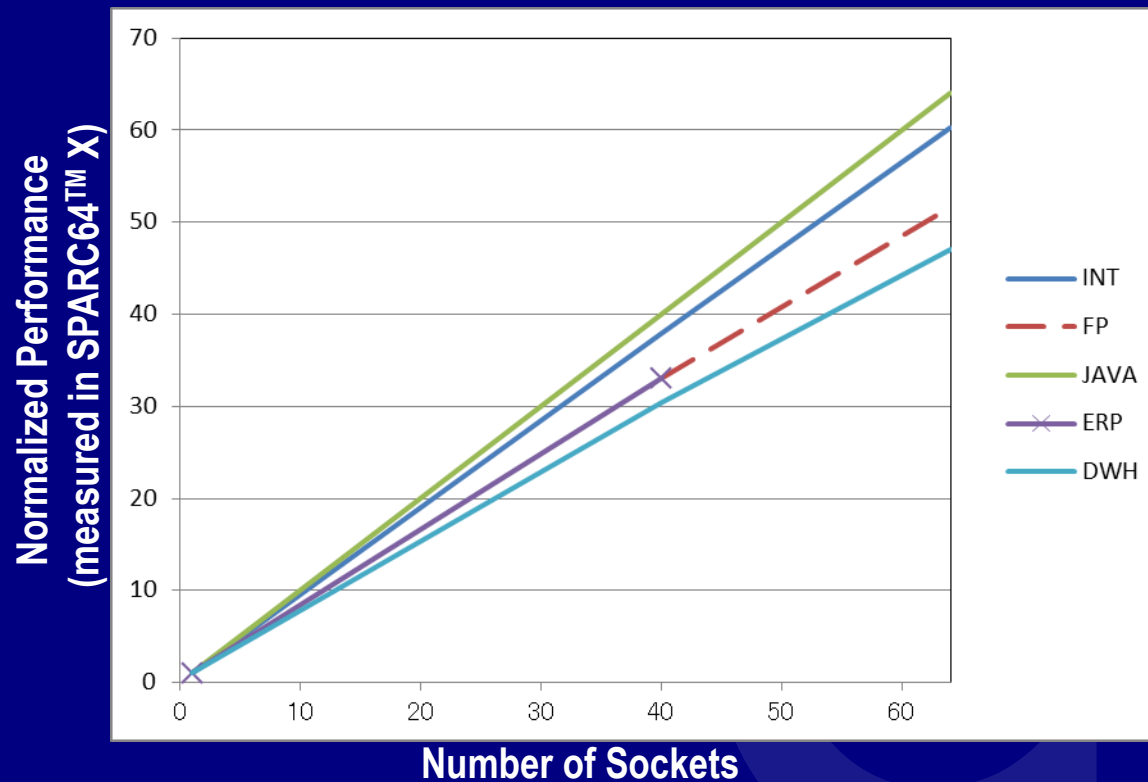
(Each line represents connections between a BB and two XBs in a XB-Box)



System Scalability

- ◆ SPARC64™ X systems demonstrate high scalability across a wide-range of applications
 - Integer, Floating-Point, Java, ERP, DWH

SPARC64™ X efficiently scales to 64 CPU sockets



Agenda

◆ Fujitsu Processor Development

◆ SPARC64™ X+

- Design Concept and Processor Overview
- Software on Chip (SWoC)
- Micro-Architecture
- System Architecture
- RAS
- Power Management

◆ Summary

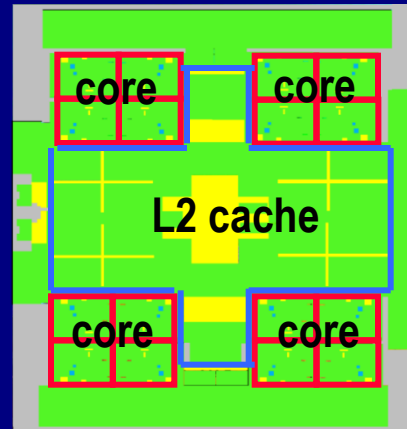


Reliability, Availability, Serviceability

Units	Error Detection and Correction
Cache (Tags)	ECC, Parity & Duplicate
Cache (Data)	ECC, Parity
Registers	ECC (INT/FP), Parity (Others)
ALUs	Parity, Residue

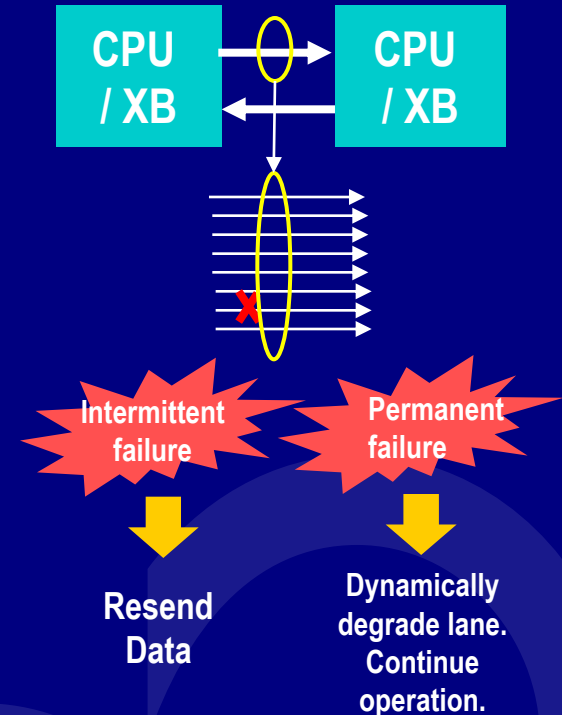
Other RAS features
Cache dynamic degradation
Hardware Instruction Retry
Lane dynamic degradation

SPARC64™ X+ RAS diagram



Green: 1-bit error Correctable
 Yellow: 1-bit error Detectable
 Gray: 1-bit error Harmless

System Bus Reliability



◆ Mainframe-level RAS features for SPARC64™ X / X+

- Number of checkers increased to ~54,000
- System bus mechanisms for self-recovery and lane dynamic degradation

→ Guarantee Data Integrity and Keep on Running

Agenda

◆ Fujitsu Processor Development

◆ SPARC64™ X+

- Design Concept and Processor Overview
- Software on Chip (SWoC)
- Micro-Architecture
- System Architecture
- RAS
- Power Management


◆ Summary



Power Management

◆ Save energy while Idle

- CPU Lower Power (LP) State introduced in SPARC64™ X
 - Dynamically decrease frequency and voltage
 - Keep all data and caches coherent
 - State transition managed by software
- ✓ 45% power savings measured in SPARC64™ X
- ✓ Transition time between states is ~1.7ms
- ✓ Continue working while in transition
- DIMM power saving mechanism
 - Memory controller supports two lower power states
 - Power-down
 - Self-refresh



Agenda

◆ Fujitsu Processor Development

◆ SPARC64™ X+

- Design Concept and Processor Overview
- Software on Chip (SWoC)
- Micro-Architecture
- System Architecture
- RAS
- Power Management

◆ Summary



Summary

- ◆ SPARC64™ X+ is Fujitsu's latest SPARC processor, designed for Fujitsu's next generation UNIX servers
- ◆ SPARC64™ X+ realizes improved single-thread performance with a higher clock speed, micro-architectural enhancements, and SWoC
- ◆ SPARC64™ X / X+ systems realize high scalability, from 1 to 64 CPU sockets (2048 threads)
- ◆ SPARC64™ X+ implements extensive RAS features
- ◆ Fujitsu will continue to develop the SPARC64™ series

Abbreviations

- SPARC64™ X+
 - RSA: Reservation Station for Address generation
 - RSE: Reservation Station for Execution
 - RSF: Reservation Station for Floating-point
 - RSBR: Reservation Station for Branch
 - GUB: General-purpose Update Buffer
 - FUB: Floating-point Update Buffer
 - GPR: General-Purpose Register
 - FPR: Floating-Point Register
 - CSE: Commit Stack Entry
 - EAG: Effective Address Generator
 - EX : Execution unit (Integer)
 - FL : Floating-point unit
 - HPC-ACE: High Performance Computing-Arithmetic Computational Extensions
 - ERP: Enterprise Resource Planning
 - DWH: Data WareHouse