SPARC64™ XII:
Fujitsu's latest 12 Core Processor for Mission Critical Servers

April 20, 2017

Takumi Maruyama

Advanced System Research & Development Unit
Fujitsu Limited

Agenda

- Fujitsu Processor Development History

- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance

- Summary

Fujitsu Processor development

Perpetual evolution

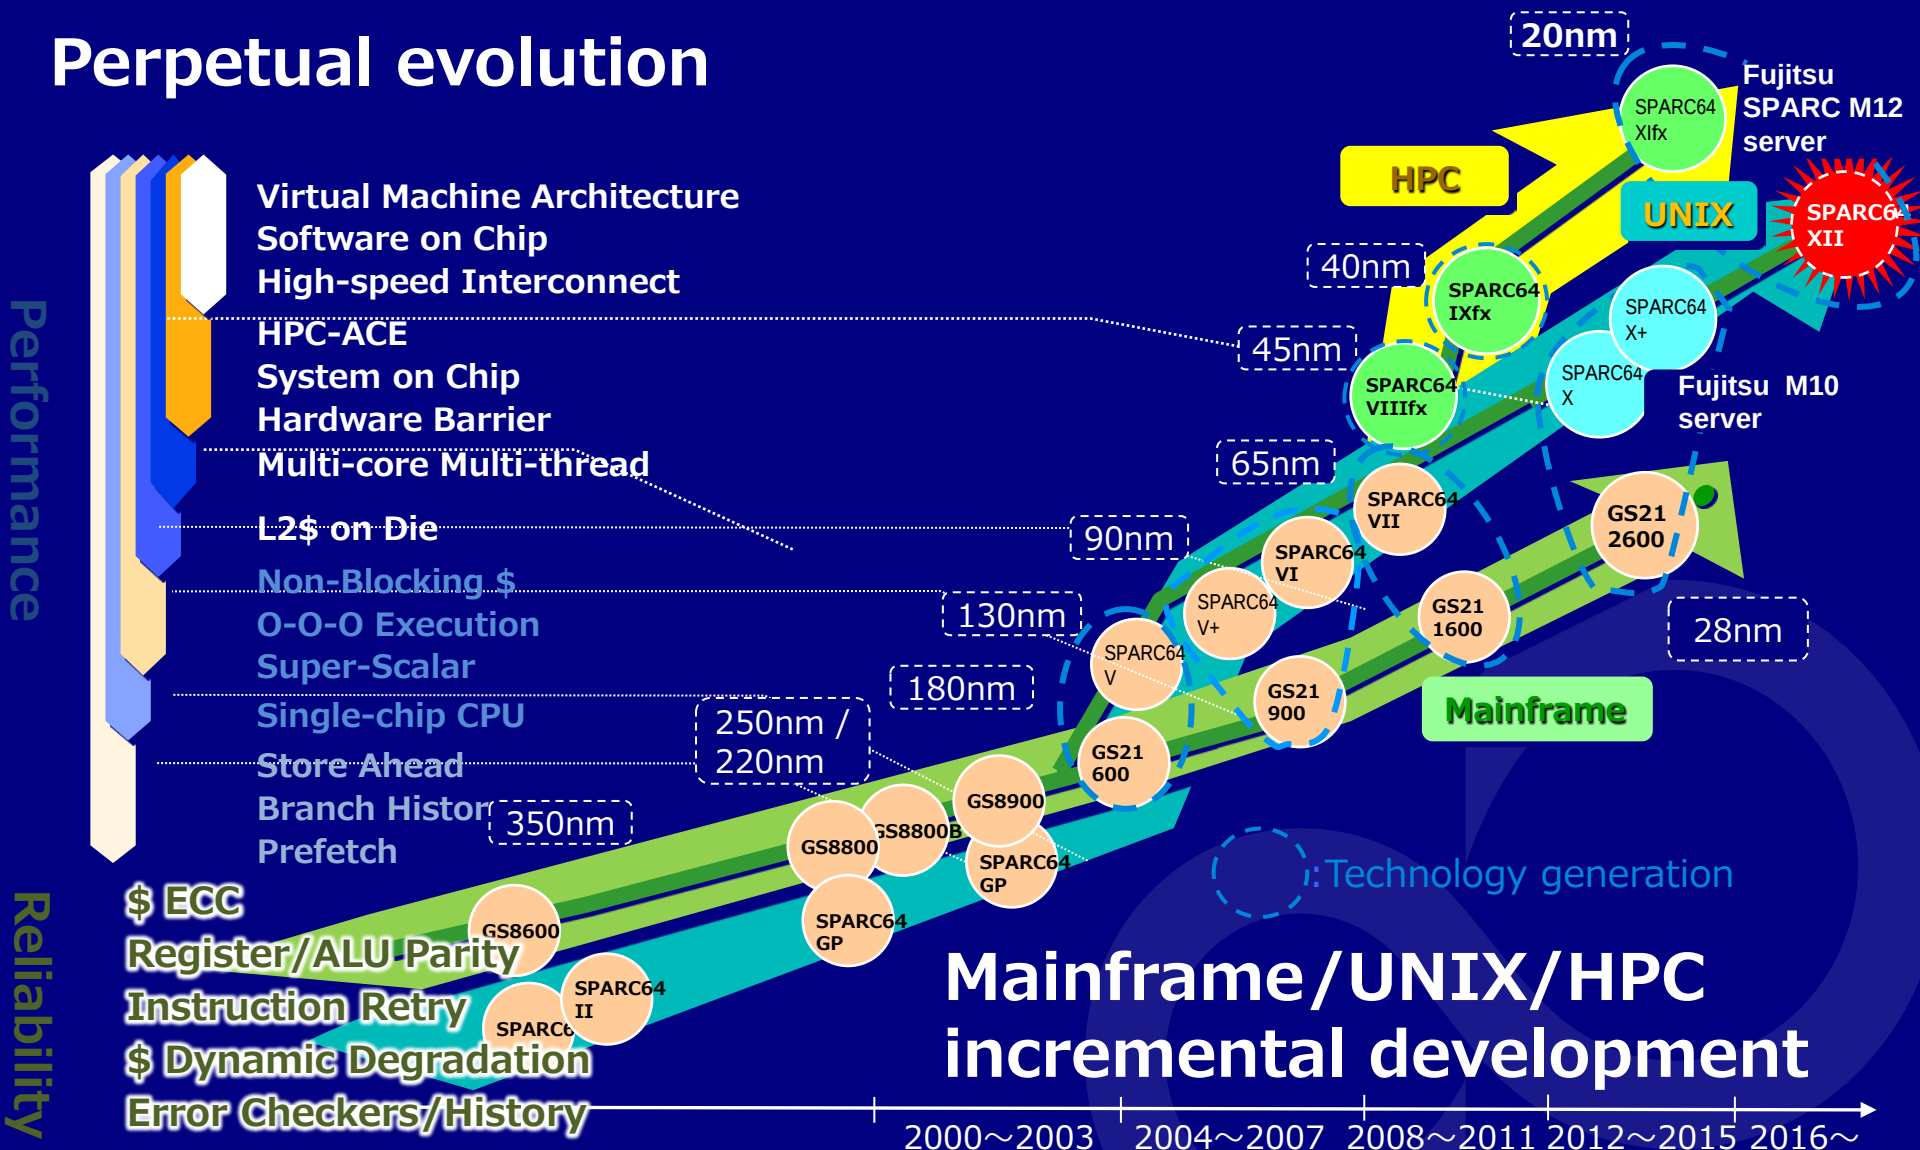

Agenda

- Fujitsu Processor Development History
- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance
- Summary

SPARC64™ XII Design Concept

High Performance: SPEED and THROUGHPUT

- **High SPEED (Single thread performance)**
  - High CPU frequency
  - State of the art O-O-O
  - Rich execution units
- **High THROUGHPUT**
  - Many cores, threads
  - Strong Cache and Memory
- **Robust RAS**

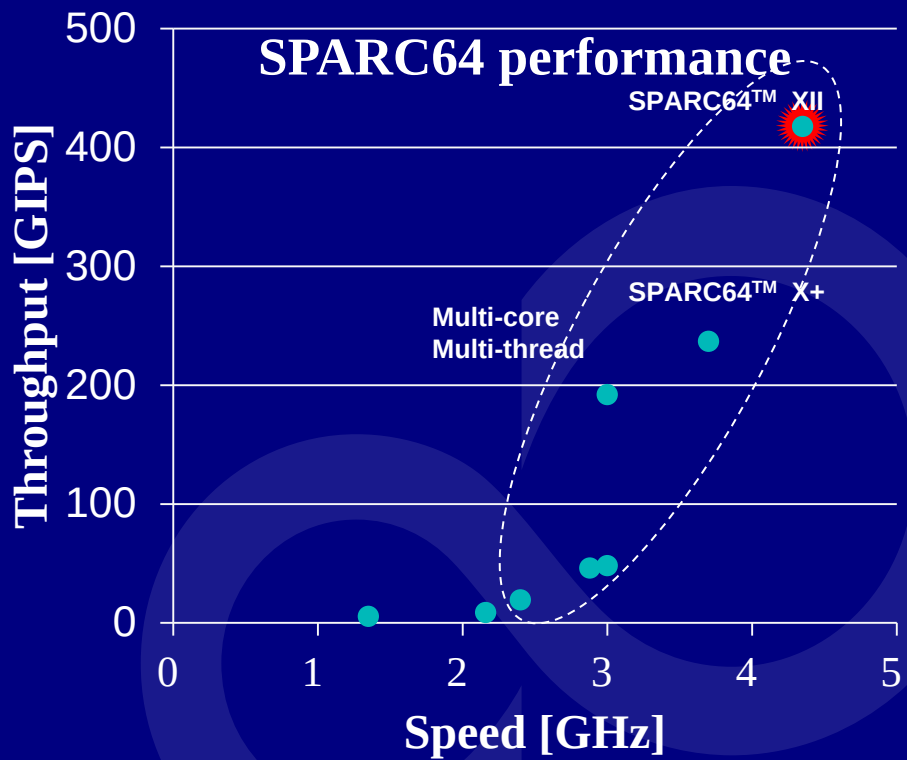

Agenda

- Fujitsu Processor Development History
- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance
- Summary

SPARC64™ XII Chip Overview

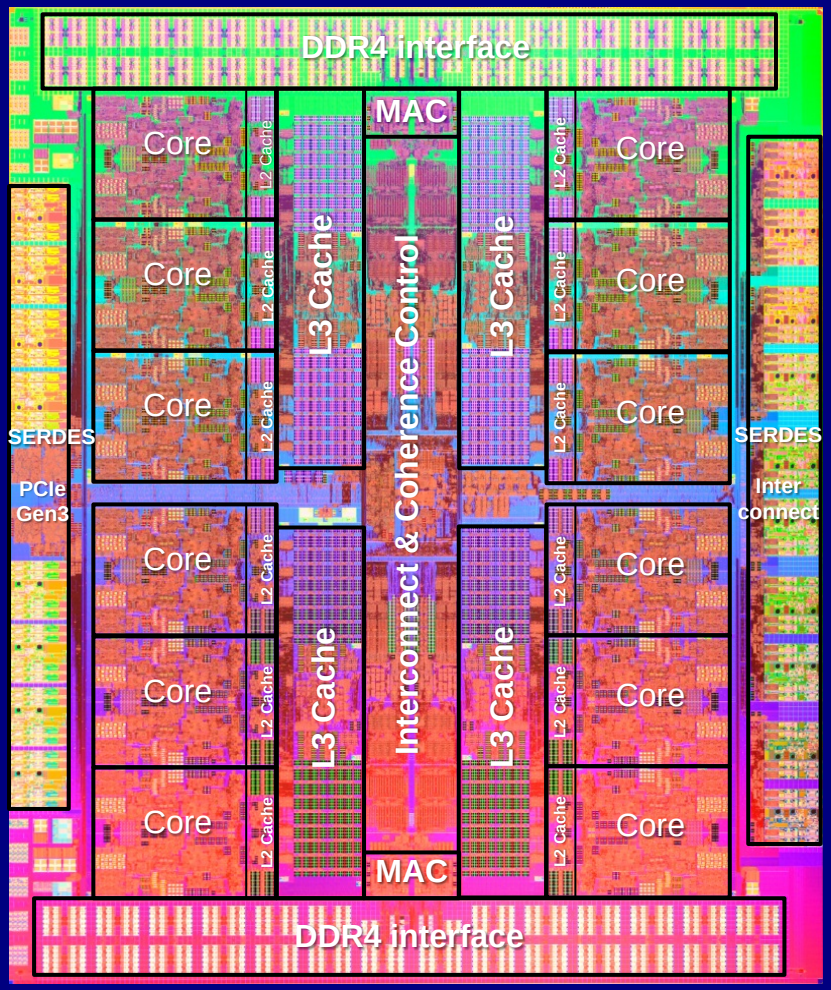

- **Architecture Features**
  - 12 cores x 8 threads
  - SWoC (Software on Chip)
  - 32MB L3 cache
  - Embedded Memory and IO Controller (PCIe GEN3)
- **20nm CMOS**
  - 25.8mm x 30.8mm
  - 5,450M transistors
  - 1,860 signal pins
  - 4.25GHz (up to 4.35GHz with High Speed Mode enabled)
- **Performance (peak)**
  - 417GIPS / 835GFlops
  - 153GB/s memory throughput

SPARC64™ XII Core spec

Instruction Set Architecture	SPARC-V9/JPS HPC-ACE VM / SWoC
Integer Execution Units	(156 GPR x 4 + 96 GUB) x 2pipe
	(ALU/SHIFT x2) x 2pipe (ALU/AGEN x2) x 2pipe (MULT/DIVIDE x1) x 2pipe
FP Execution Units	(128 FPR x 4 + 64 FUB) x 2pipe
	(FMA x4 / FDIV x2) x 2pipe (IMA/Logic x4) x2pipe (Decimal x1 / Cypher x2) x 2pipe
L1$	L1I$ 64KB/4way L1D$ 32KB/4way x 2pipe
L2$	512KB/16way

SPARC64™ XII Pipeline

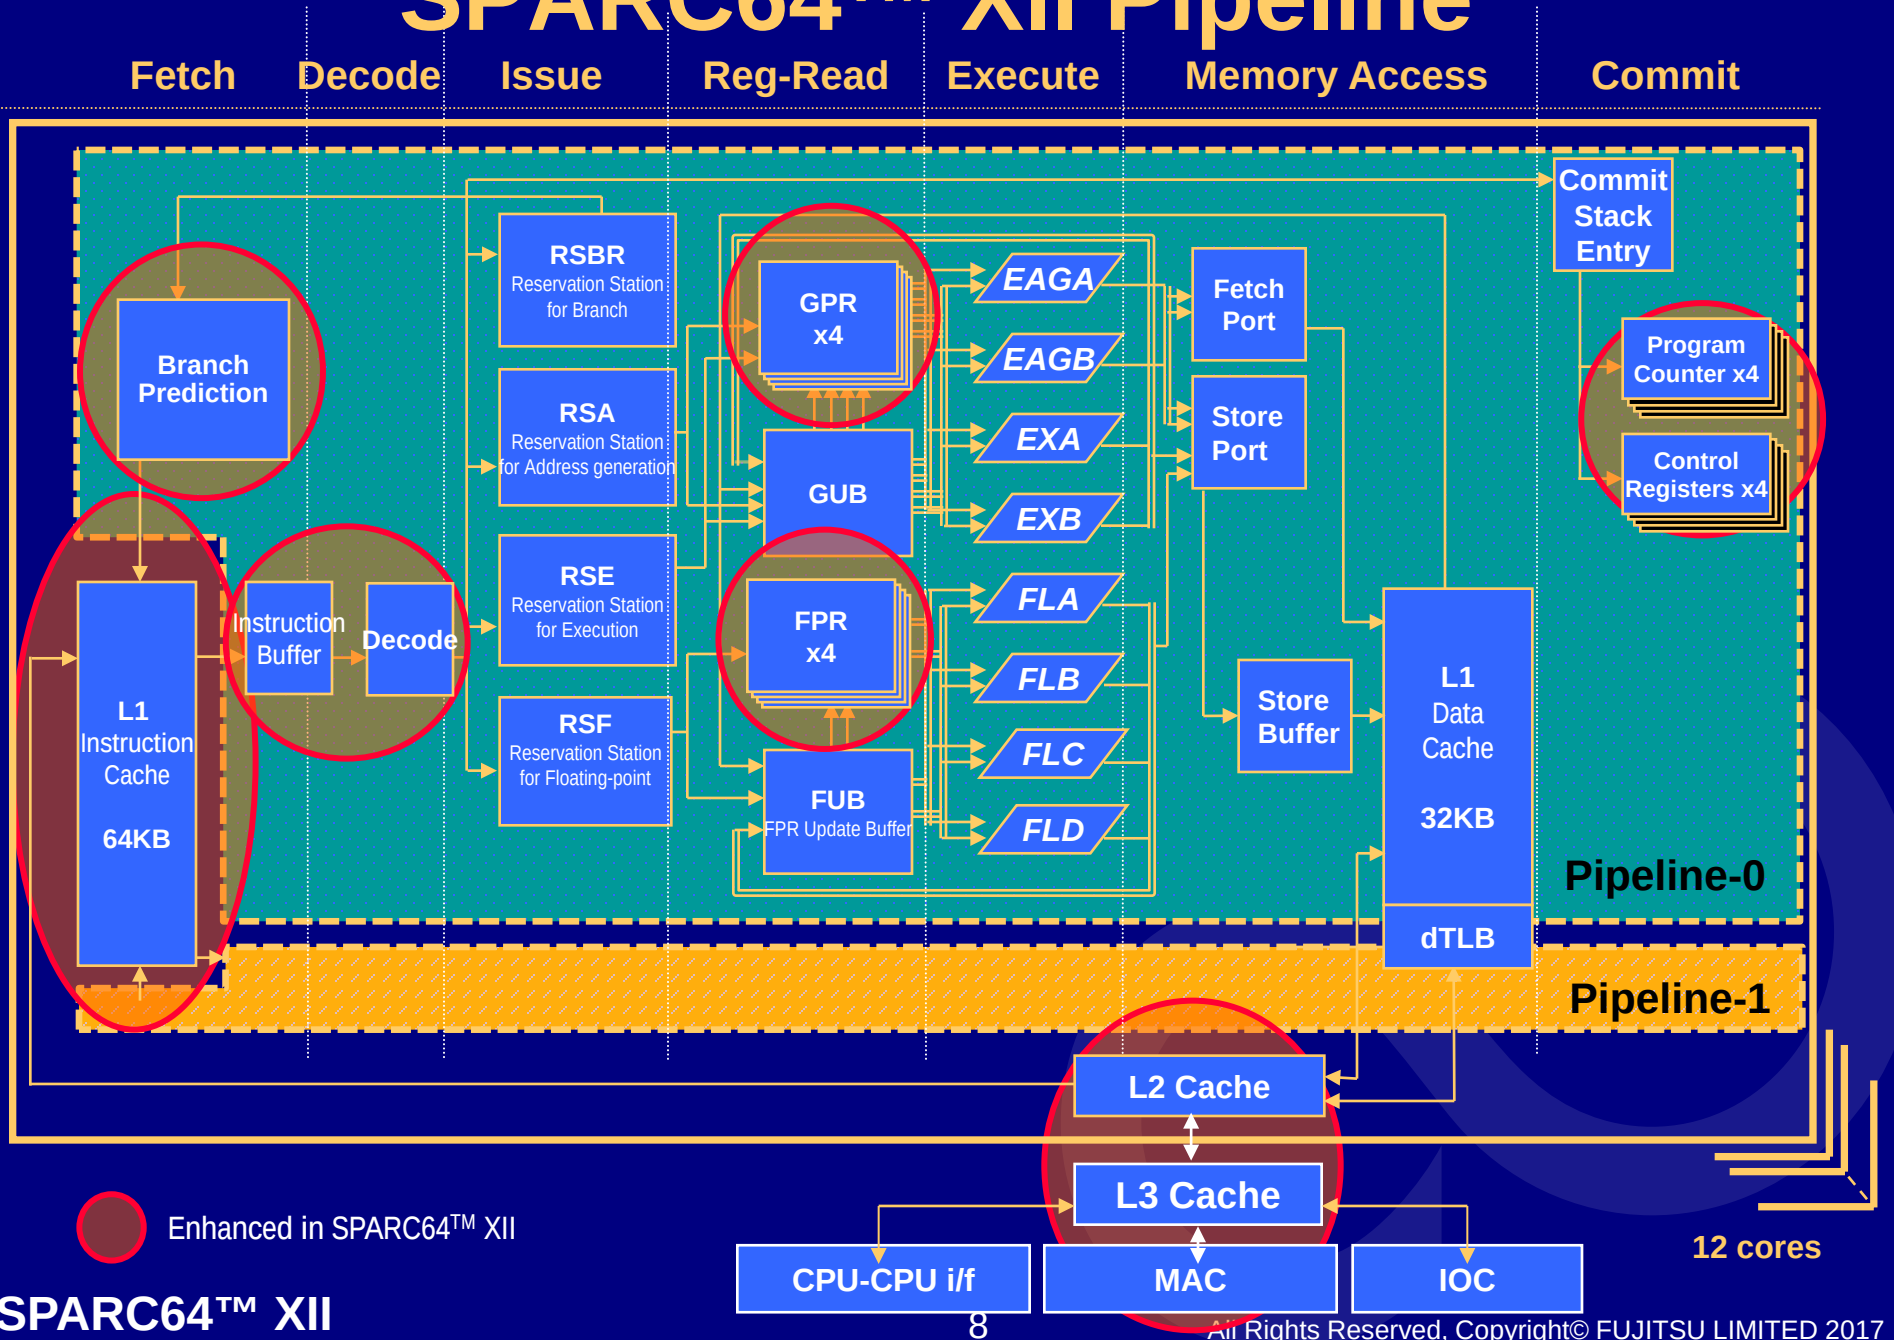

Agenda

- Fujitsu Processor Development History
- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance
- Summary

Common Micro-Architecture features between SPARC64™ X+ and XII

- **Processor Core**
  - Fetch 8 instructions, Decode 4 instructions, Execute 6 instructions per instruction pipeline
  - Aggressive O-O-O execution, including memory access (load and store)
  - High single thread performance
  - SWoC (Software on Chip) Accelerates specific software function by hardware
- **Outside the Core**
  - Embedded Memory and IO Controller (PCIe GEN3)
  - Robust RAS (Reliability, Availability, Serviceability) ★

Reliability, Availability, Serviceability

Units	Error detection and correction scheme
Cache (Tag)	ECC Duplicate & Parity
Cache (Data)	ECC Parity
Register	ECC (INT/FP) Parity(Others)
Execution Unit	Parity/Residue
Cache dynamic degradation	Yes
HW Instruction Retry	Yes
History	Yes

→ Guarantees Data Integrity
(Number of checkers increased to ~66,000)

SPARC64™ XII RAS diagram

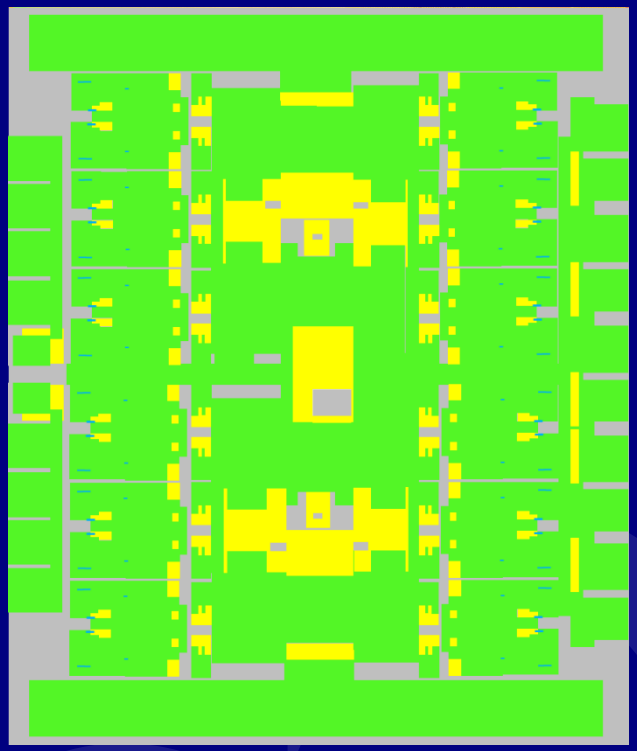

Green: 1bit error Correctable
Yellow: 1bit error Detectable
Gray: 1bit error harmless

Hardware Instruction Retry

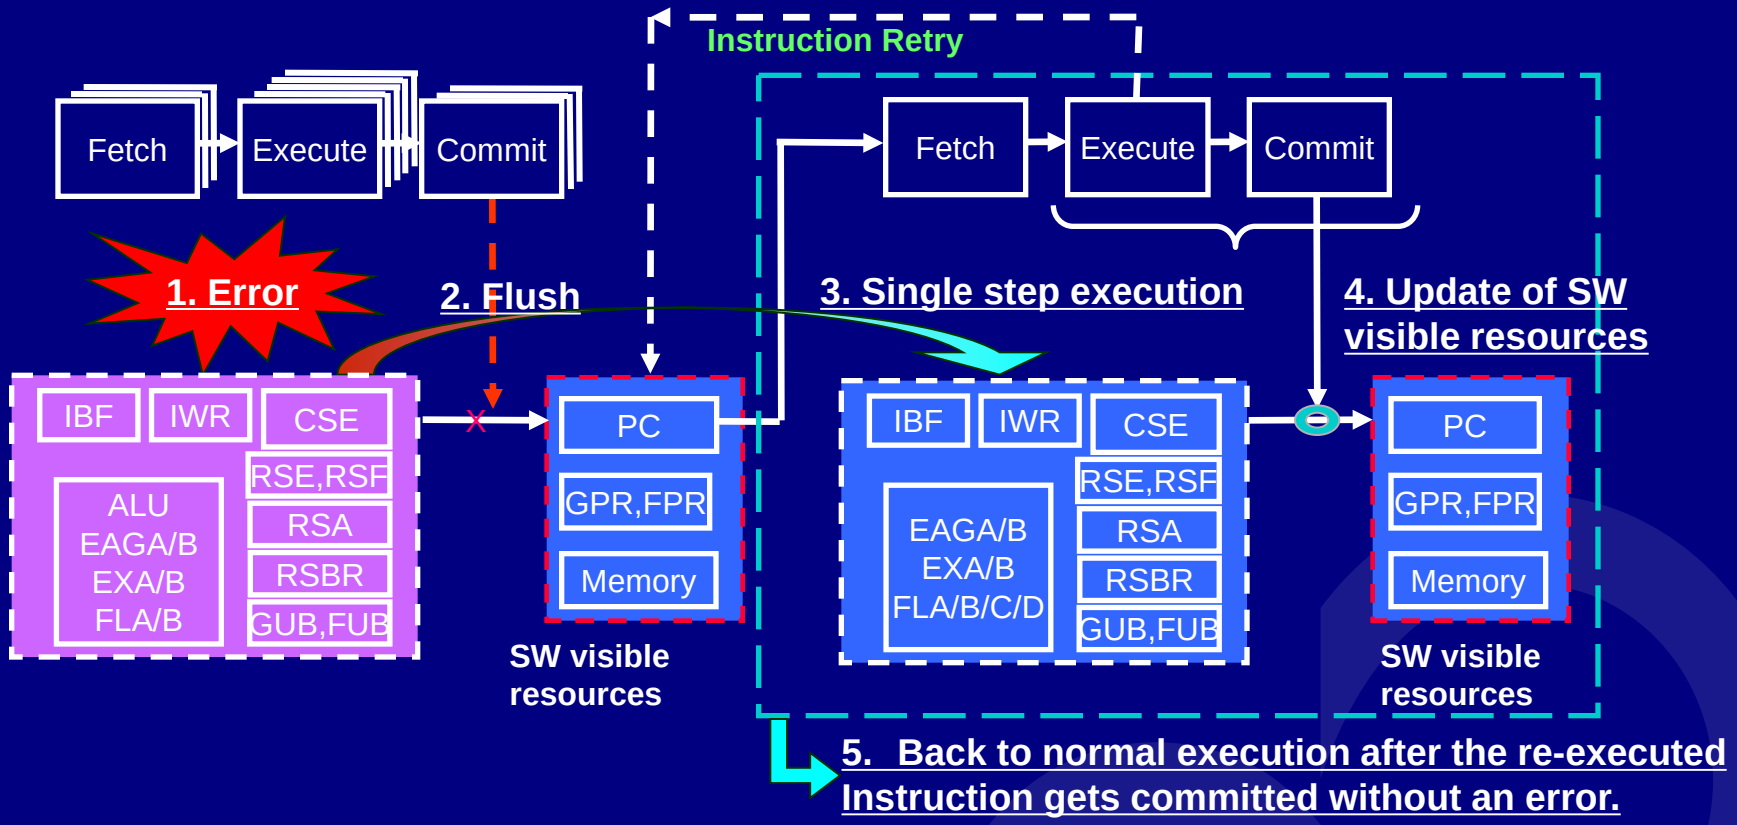

- When an error is detected, Hardware re-execute the instruction automatically to remove the transient error by itself.

Enhanced Micro-Architecture from SPARC64™ X+

- **Processor Core**
  - Dual instruction pipelines sharing L1 I-cache, and 8-way SMT (4-way SMT per instruction pipeline) ★
  - Better Branch Prediction Scheme
  - Various Queue-size increase
  - Deeper pipeline to increase Frequency
  - High Speed Mode
- **Outside the core**
  - 3 level hierarchy cache, and 4 LCU (Last level cache and Core Unit) structure ★
  - DDR4-2400 memory
  - Doubled PCIe GEN3 ports (8 lane x 4)

Micro-architecture enhancements 1/2

Increase core throughput, keep single thread performance

- Dual instruction pipelines per core
  - Each pipeline has its own resources except ↓
- Sharing L1 I-Cache and TLB between the two pipelines
  - Especially effective in DB apps
- 4-way SMT instruction pipeline
  - More throughput compared with the previous 2-way SMT
  - Resources are dynamically shared between the threads.

Core Resource allocation

#Active threads	L1 I$	IB	Rsv. Station	Rename Registers	Execution Units	FP/SP Port	L1 D$	CSE
1	100%	50%	50%	50%	50%	50%	50%	50%
2	100%	50% x2	50% x2	50% x2	100%	50% x2	100%	50% x2
4	100%	25% x4	25% x4	25% x4	100%	25% x4	100%	25% x4
8	100%	12.5% x8	50% x2	12.5% x8	100%	12.5% x8	100%	12.5% x8

Micro-architecture enhancements 2/2

3 level cache and 4 LCU structure to realize Low Latency and High throughput

- L2 cache latency < L3 cache latency x 0.5
- L2 cache to Core : 32 Bytes x 2 pipelines
- L3 cache to Core : 32 Bytes x 3 cores

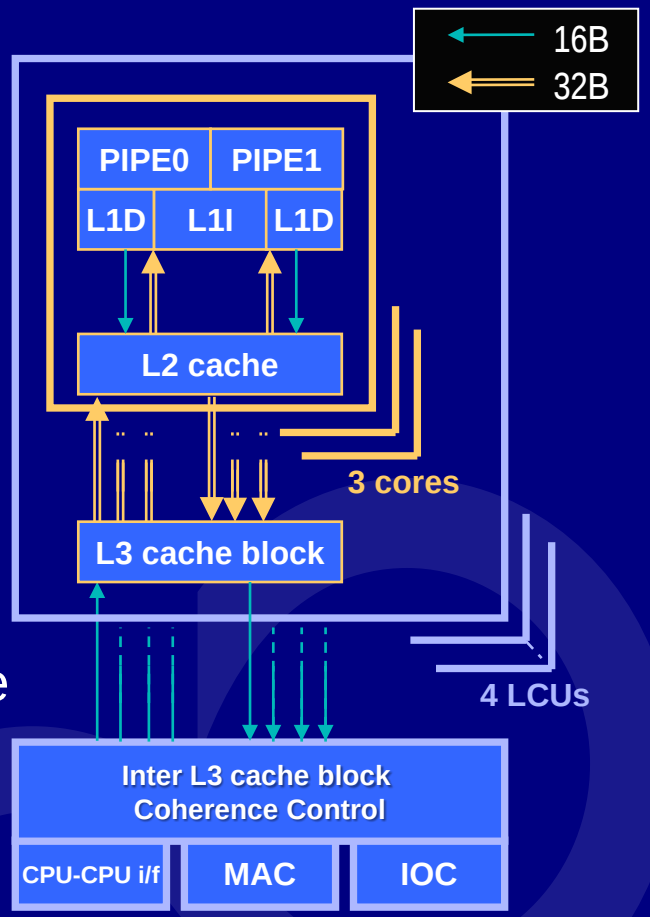

- MEZASI
  - Fujitsu's unique cache protocol for LCU
- ELC (Extra Level Cache)
  - Unused L3 cache blocks act as Victim Cache
- IO-DCA (IO Direct Cache Access)
  - Direct DMA write to L3 cache from IO device
- Speculative memory access
  - Memory access in parallel with other L3 cache blocks access

Agenda

- Fujitsu Processor Development History
- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance
- Summary

Fujitsu SPARC M12 server

- Fujitsu's latest UNIX server with SPARC64™ XII.
  - M12-2: mid-range server
  - M12-2S: scalable high-end server

M12-2S

M12-2

SPARC M12 System Architecture

- Scales from 1 to 32 CPU sockets (2,048→3,072 threads)
  - Directory-based cache coherency
  - High-speed interconnect, up to 25 Gbps per lane in SPARC64™ XII

- System Configuration
  - Building Block (BB) is 2 CPUs and 2 XBs
  - Up to 4 BBs can be connected by XBs
  - 16 BBs can be connected via XB-Boxes
  - BB topology inherits current M10-4S with SPARC64™ X+

Building Block (2 CPU Sockets)

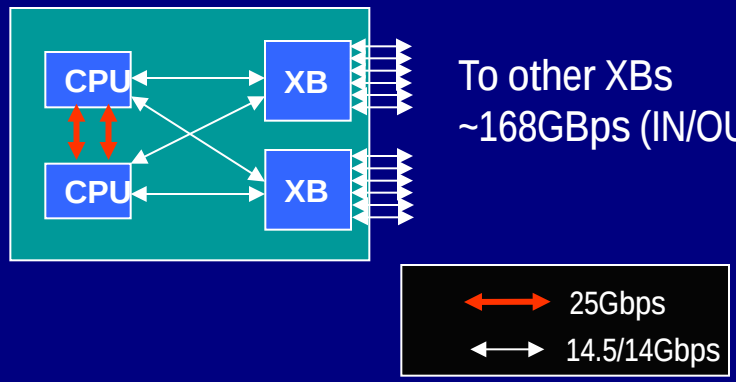

16 BBs (32 CPU Sockets)

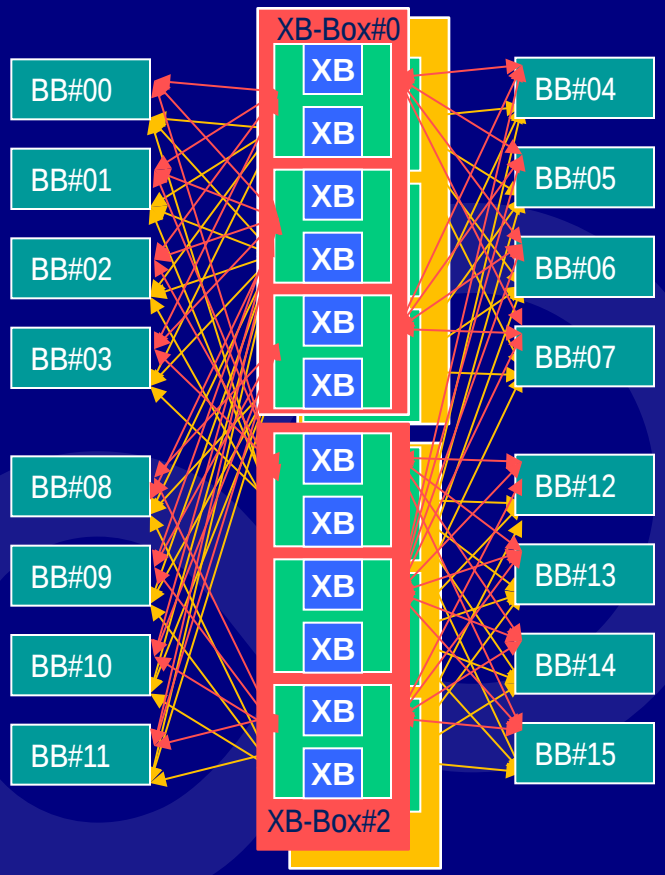

SPARC M12 System Cooling

- VLLC (Vapor and Liquid Loop Cooling)
  Newly developed cooling mechanism for SPARC M12.
  - Evaporative cooling, taking heat away when liquid changes to vapor
  - The pumps circulate coolant in the VLLC system
  - The radiator dissipates the heat absorbed by the cooling plate into the air.

➔ Achieves 2x cooling performance of the current LLC

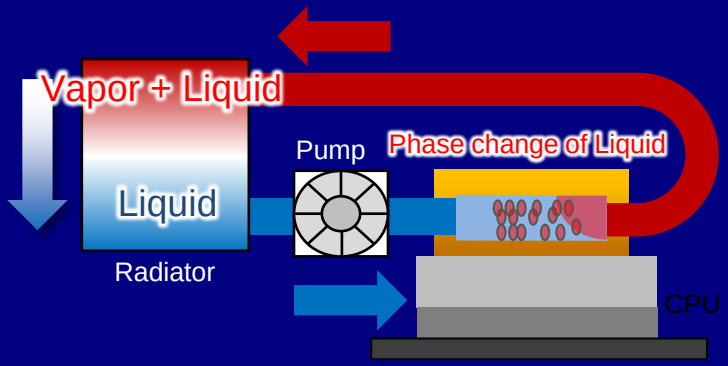

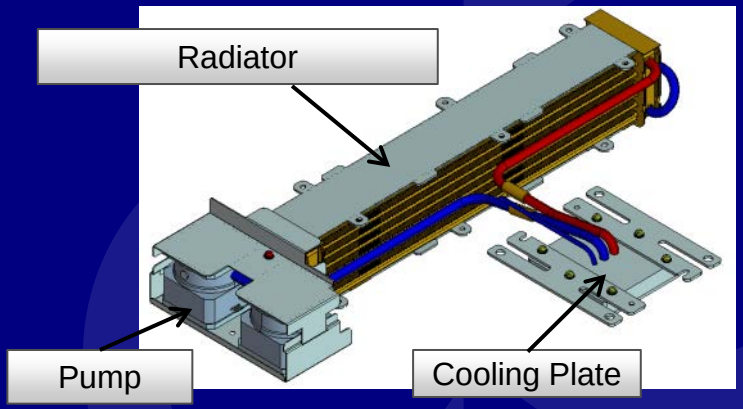

Agenda

- Fujitsu Processor Development History
- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance
- Summary

SPARC64™ XII benchmark numbers

		1CPU	2CPU	16CPU
SPECint®_rate2006		956	1,910	14,500
SPECjbb®2015 Composite	Max Critical	52,659 jOPS 34,771 jOPS	98,333 jOPS 63,354 jOPS	n/a n/a
SPECjbb®2015 Multiple-JVM	Max Critical	54,434 jOPS 34,771 jOPS	n/a n/a	n/a n/a
STREAMtriad		127GB/s	n/a	1,533GB/s

Source except STREAMtriad : www.spec.org

Fujitsu SPARC M12-2S：
SPARC64 XII（4.25GHz, up to 4.35GHz）x 1,2,16CPU
OS：Oracle Solaris 11.3
Compiler: Oracle Developer Studio 12.6
JVM：HotSpot™ 64-Bit Server VM, version 1.8.0_121

SPARC64™ XII Performance

- **2x chip throughput and 2.5x core throughput**
  compared to previous SPARC64™ XII, keeping single thread performance
- **Keys: 8way SMT design, Higher CPU frequency, and micro-architecture improvement**

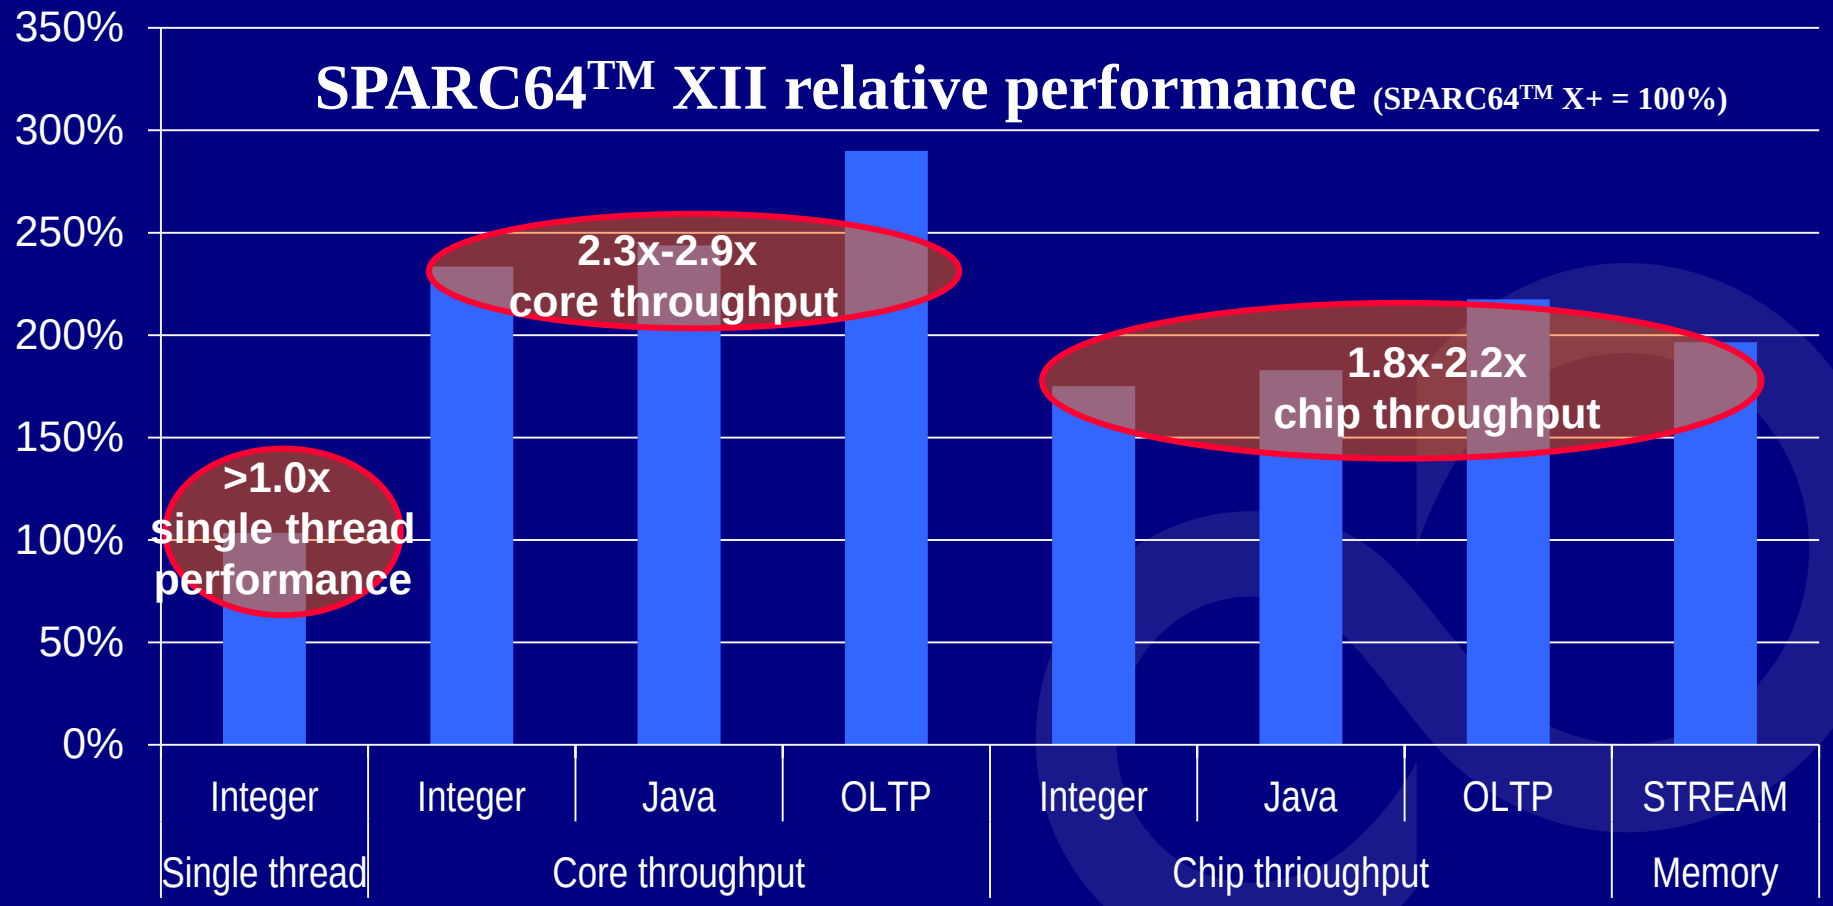

Agenda

- Fujitsu Processor Development History
- SPARC64™ XII
  - Design concept
  - Chip overview
  - Micro Architecture
  - System Architecture
  - Performance
- Summary

Summary

- SPARC64™ XII is Fujitsu's 12th SPARC processor which has been designed to be used for Fujitsu's latest UNIX server.
- The processor enhancements are to increase throughput, while keeping its high single thread performance and its robust RAS features
- SPARC64™ XII measured results have shown 2x chip throughput of SPARC64™ X+
- Fujitsu will continue to develop SPARC64™ series to meet the needs of a new era.