Hardware Overview

SCICOMP

IBM
July, 2006

Agenda

- Hardware
- Software
- Documentation

Hardware Overview

- Core:

- Processors:

- Nodes:

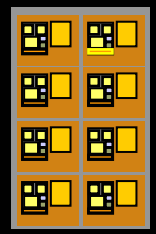

- Clusters:

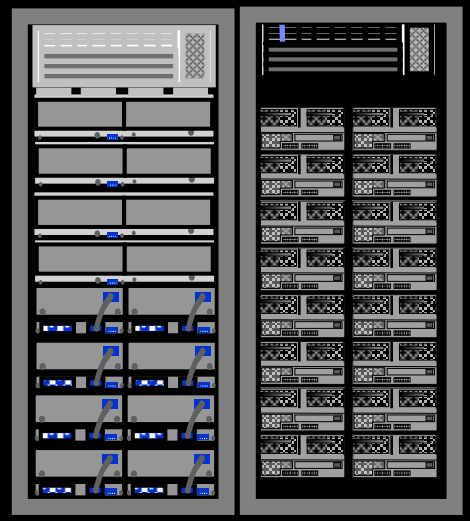

IBM Product Naming

New Name	Old Names	Market	Processor
System i	iSeries, AS400	Commercial	RS64 POWER5
System p	RS6000 SP pSeries	Server, technical	POWER3 POWER4 POWER5
System x	xSeries IA-32	Server, technical	Intel AMD PowerPC
System z	zSeries ES9000	Mainframe	zSeries

Power Processor Progression

Processor	Years	Clock Rate	Feature
POWER2	1990 - 1994	20 – 60 MHz	RISC
P2SC	1994 - 1998	60 – 150 MHz	Bandwidth
POWER3	1998 – 2002	200 – 450 MHz	Single Chip
POWER4	2001 – 2005	1 – 1.9 GHz	Dual Core
POWER5	2004 -	1.5 – 1.9 GHz	Multi-Thread
POWER5+	2006-	1.9 – 2.2 GHz	Speed bump

POWER5 Systems

- POWER5 processors
  - Single and Dual processor chips
- Modules
  - Dual Chip Modules (DCM)
  - Multi Chip Modules (MCM)
- Nodes
  - Multiple modules
    - p5-575
    - p5-595
  - SMP within a node
- Cluster
  - Multiple nodes
  - Connected with High Speed Switch (HPS)

System p5 “Nodes” – partial list

Model	Processors	Clock Rate (GHz)	Max Memory (x 2^30 byte)
p5 595	16-64	1.65, 1.9	2000
p5 590	8-32	1.65	1000
p5 575	8-16	1.9, 2.2*	256
p5 570	2-16	1.9, 2.2*	512
p5 560Q	4-16	1.5*	128
p5 520	1,2	1.65, 1.9*	32
p5 505	1,2	1.5, 1.65*	32

* - POWER5+

POWER5 Processor Systems

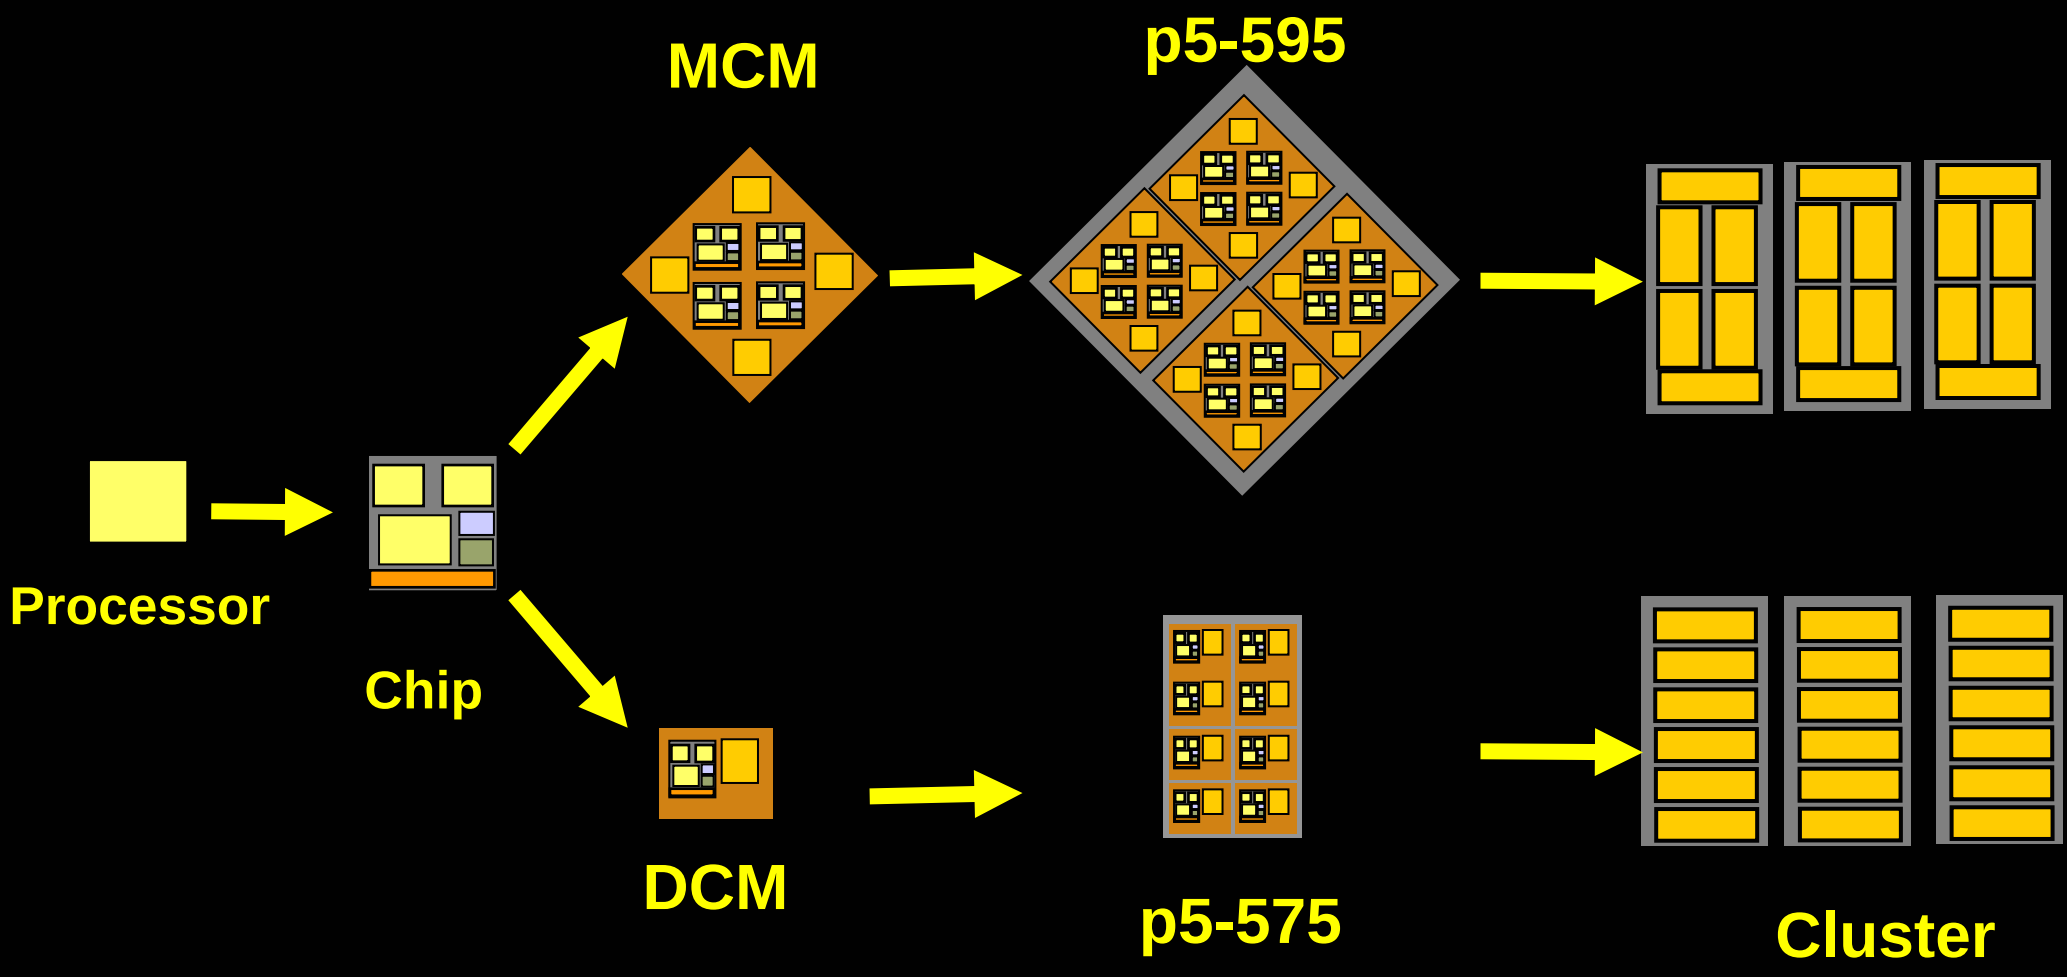

POWER5 Features

- Private L1 cache
- Shared L2 cache
- Shared L3 cache
- Interleaved memory
- Hardware Prefetch
- Multiple Page Size support

Processor Characteristics

- High frequency clocks
  - Deep pipelines
  - High asymptotic rates
- Superscalar
- Speculative out-of-order instructions
- Up to 8 outstanding cache line misses
- Large number of instructions in flight
- Branch prediction
- Hardware Prefetching

Block Diagram

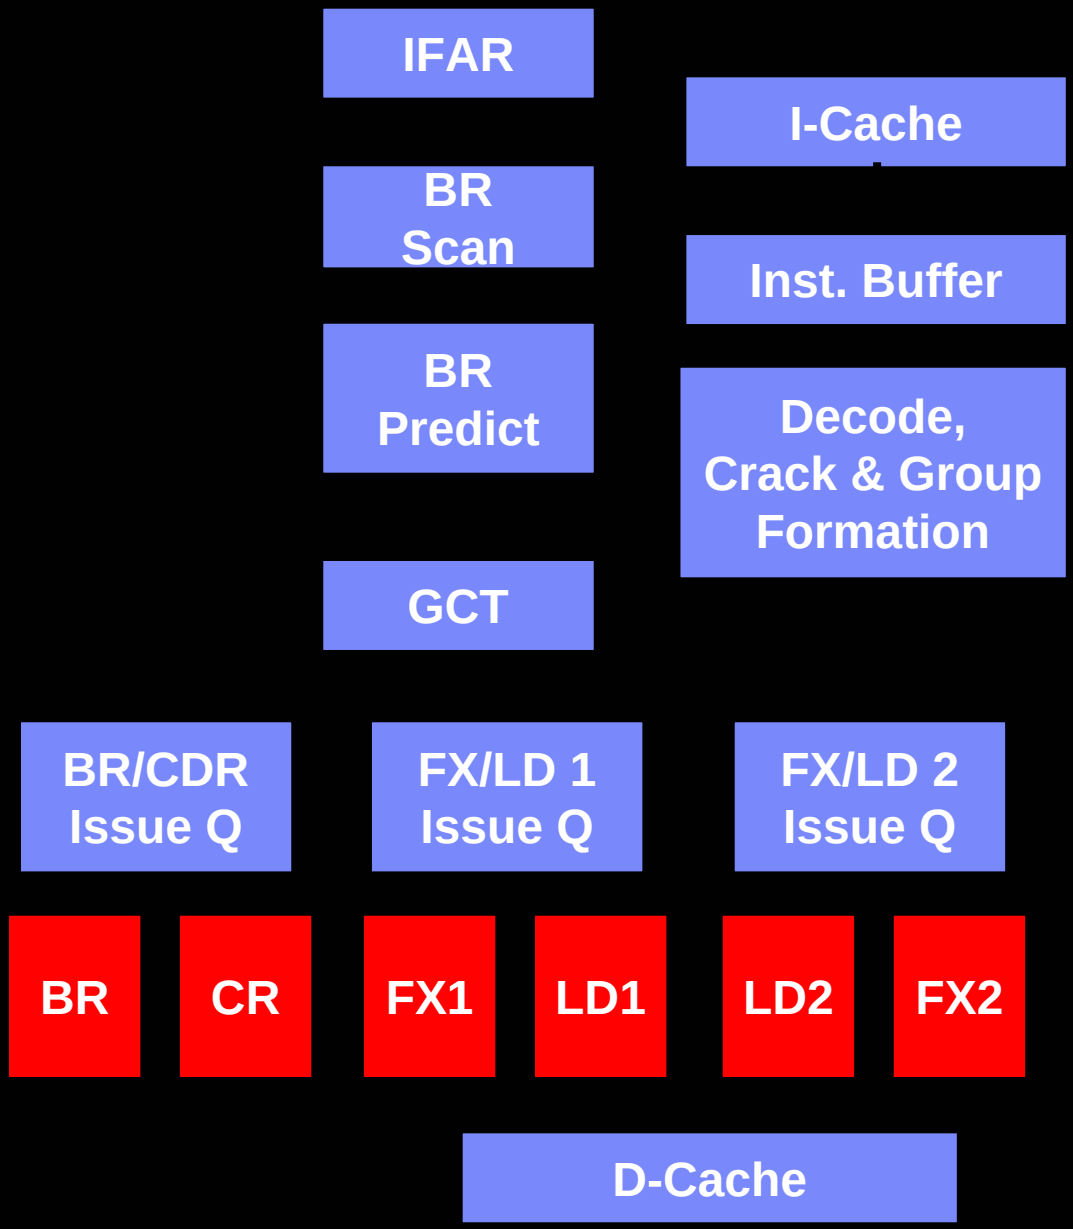

- Eight-way Issue
- Eight independent functional units

Processor Features

	POWER4	POWER5
Clock	1.0 – 1.9 GHz	1.5 – 2.2 GHz
Caches	Three levels	Three levels
L3 Speed	1/3 clock frequency	½ clock frequency
Virtualization	Up to 32 partitions	Up to 254 partitions
Partitions	Unit processor	Fractional
Power Mang.	Static	Dynamic
Thread Execution	Single Thread	Multi Threading
Memory Store	Single Buffer	Double Buffer
Renaming Registers	GP: 72 FP: 80	GP: 120 FP: 120

Caches and Memory

	POWER4	POWER5
L1 Cache	Data: 32 kbyte Instruction: 64 kbyte 2-way Assoc., FIFO	Data: 32 kbyte Instruction: 64 kbyte 4-way Assoc., LRU
L2 Cache	1.5 Mbyte 8-way Assoc., FIFO	1.9 Mbyte 10-way Assoc., LRU
L3 Cache	32 Mbyte 8-way Assoc., LRU 120 Cycles	36 Mbyte 12-way Assoc., LRU ~80 Cycles
Memory Bandwidth	4 Gbyte/s/Chip*	16 Gbyte/s/Chip*

* - if all memory DIMM slots occupied

POWER4 – POWER5 Comparison

	POWER4+	POWER5
Frequency (GHz)	1.7	1.9
L2 Latency (Cycles)	12	12
L3 Latency (Cycles)	120	80
Memory Latency (Cycles)	351	220
Copy Bandwidth 4 proc. (Gbyte/s)	8	18
Linpack Rate N=1000 (Gflop/s)	3.9	5.6
SPECint_base2000	1077	1398
SPECfp_base2000	1598	2576

POWER5 Design: Summary

- More gates
  - 170 million → 260 million
- Enhancements
  - Increased cache associativity
  - Increased number of rename registers
  - Reduced L3 and cache latency
- New features
  - Simultaneous Multi Threading
  - Dynamic power management

POWER5 Dual Chip Module

- One POWER5 chip
  - Single or Dual Core
- One L3 cache chip

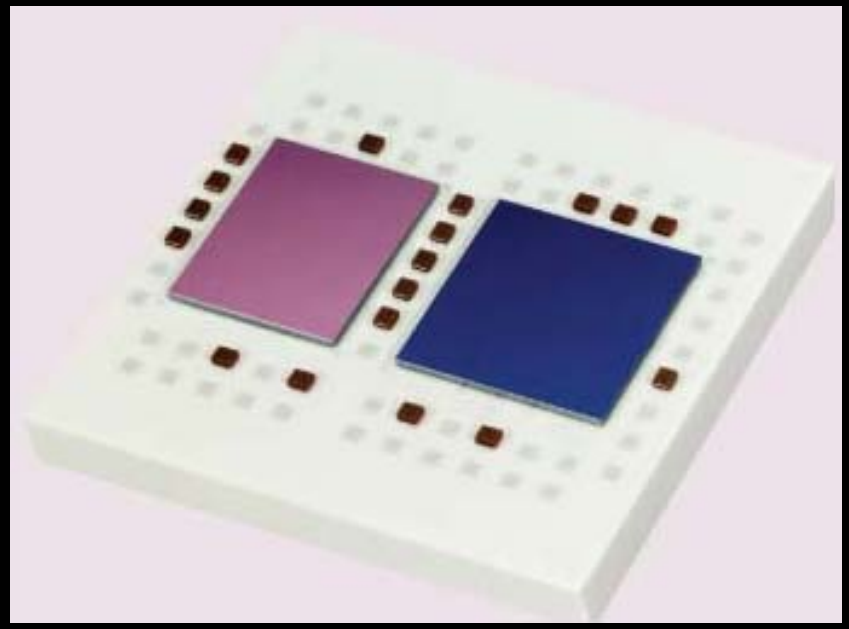

Modifications to POWER4 System Structure

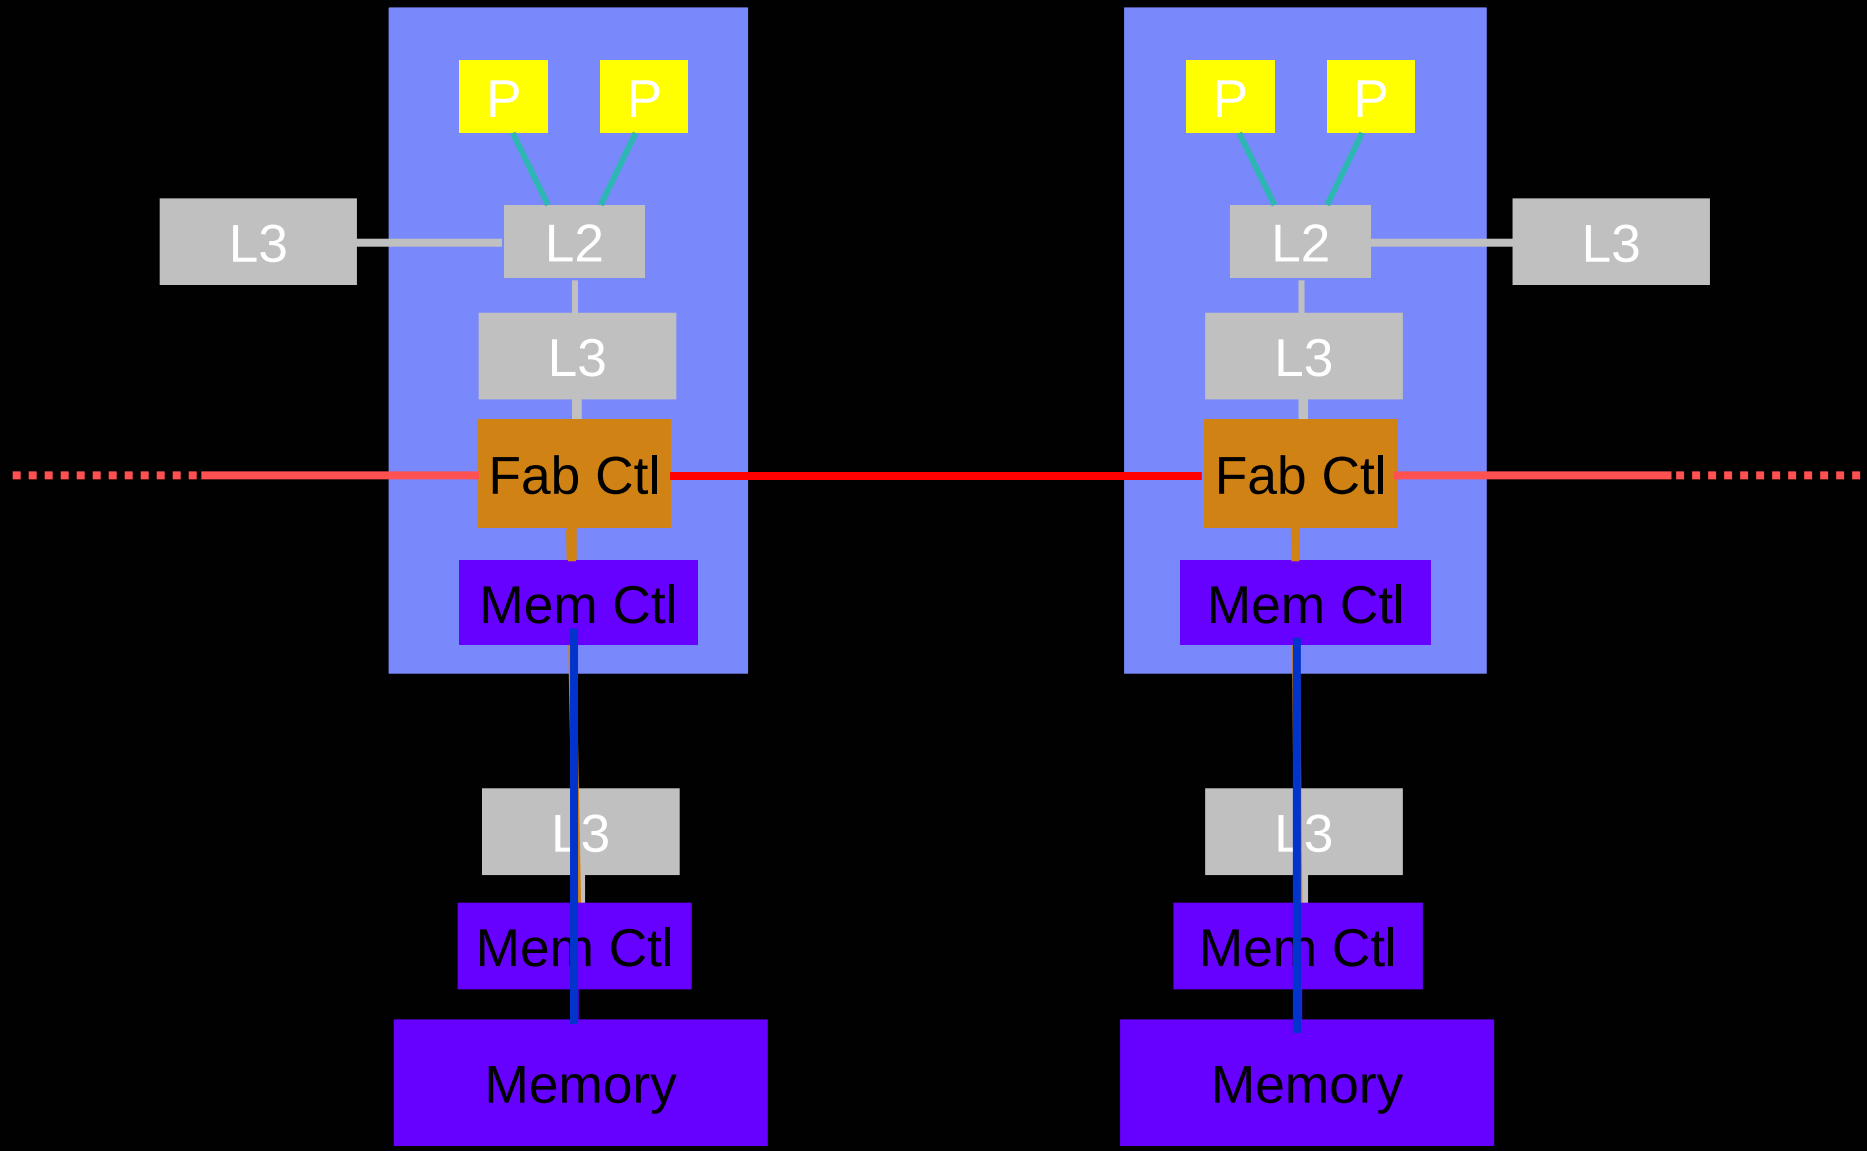

End of hardware overview