

MIPS P-Class P6600 Multiprocessor Core

The MIPS P6600 is a 64-bit processor core that represents an evolution of the MIPS P-class family.

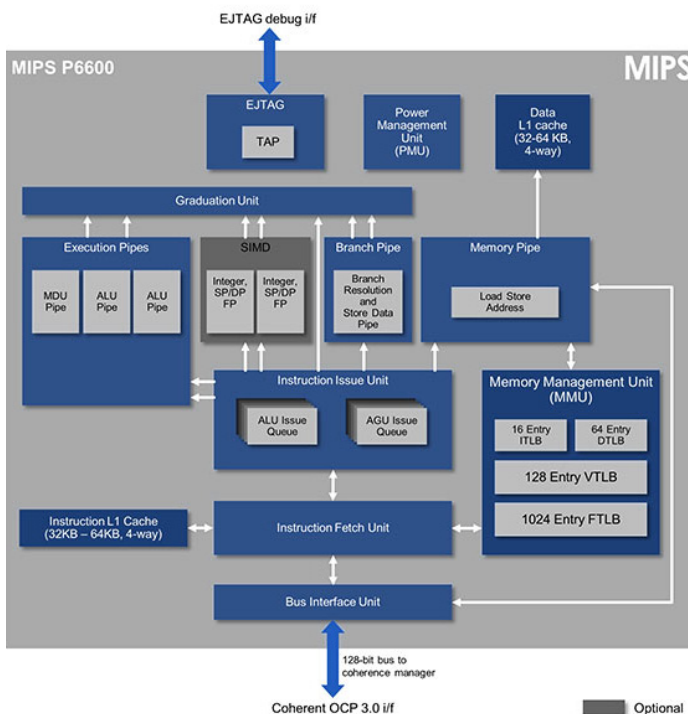
Building on the 32-bit P5600 CPU, and paving the way to future generations of high performance 64-bit MIPS processors, the P6600 is the most efficient mainstream high-performance CPU choice, enabling powerful multicore 64-bit SoCs with optimal area efficiency for applications in segments including home entertainment, networking, automotive, embedded high-performance compute and more.

The MIPS P6600 CPU is based on a wide issue, deeply out-of-order (OoO) implementation utilizing the latest release 6 of the MIPS64 architecture, supporting up to six cores in a single cluster with high performance cache coherency. Complementing this raw horsepower, the core includes 128-bit integer and floating point SIMD processing, hardware virtualization, and larger physical and virtual addressing space coming from the MIPS64 architecture.

The P6600 processor delivers performance in a smaller silicon footprint than leading IP core alternatives. SoC designers can use this efficiency advantage for cost savings, or to implement additional cores to deliver a performance advantage against competing silicon.

P6600 Benefits

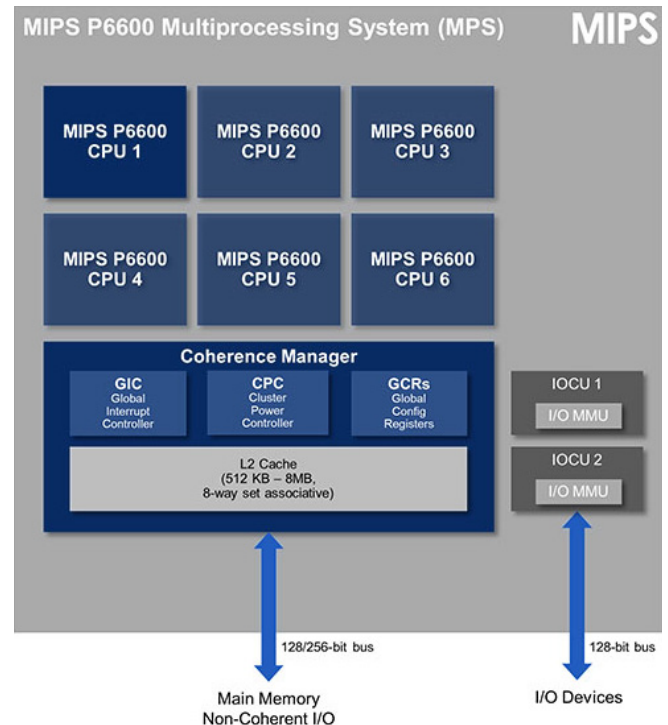
- MIPS64 r6 architecture – provided larger virtual and physical addressing, plus higher performance on 64-bit operations and data movement. Leverages latest release 6 of MIPS64, with optimizations for running JITs, Javascript, Browsers, PIC, etc.
- 128-bit SIMD – accelerates execution of audio, video, graphics, imaging, speech and other DSP-oriented software algorithms, with instruction set designed for development in high level languages such as C, OpenCL
- MIPS multi-domain security technology based on hardware virtualization – ensuring that applications that need to be secure are effectively and reliably isolated from each other, as well as protected from non-secure applications
- Multiple context security platform for enterprise/ consumer partitioning, secure content access, payments/transactions, and isolating secure schemes from numerous content sources
- Sophisticated branch prediction for maximizing utilization and performance on deeply pipelined CPU
- Load/Store bonding for optimum data movement performance
- Broad software and ecosystem support and mature toolchain
- Available as synthesizable IP for implementation in any process node, with standard cells and memories



Base Core Features

- 64-bit MIPS64® Release 6 Instruction Set Architecture
- High-performance, 16-stage, wide issue, out-of-order (OoO) pipeline
 - Quad instruction fetch per cycle
 - Triple bonded dispatch per cycle
 - Instruction peak issue of 4 integer and 2 SIMD operations per cycle
 - Sophisticated branch prediction scheme, plus L0/L1/L2 branch target buffers (BTBs), Return Prediction Stack (RPS), Jump Register Cache (JRC)
 - Instruction bonding – merges two 32-bit integer accesses into one 64-bit access, or two 64-bit integer or floating point accesses into one 128-bit access for up to 2x increase on memory-intensive data movement routines
- L1 cache size for Instruction and Data of 32KB or 64KB each, 4-way set associative
- New high-performance dual-issue 128-bit SIMD Unit – optional
 - 32 x 128-bit register set, 128-bit loads/stores to/from SIMD unit
 - Native data types:
 - 8-/16-/32-bit integer and fixed point, 16-/32-/64-bit floating point
 - IEEE-754 2008 compliant
 - Runs at full speed with CPU core
- Full hardware virtualization
 - Provides root and guest privilege levels for kernel and user space
 - Supports multiple guests, with full virtual CPU per guest = guest OSs run unmodified
 - Separate TLBs, COP0 contexts for root and guests → full isolation, fast context switching, exception and interrupt handling by root
 - HW table walk support in TLB for optimal performance
 - Complete SoC virtualization support (IOMMU and interrupt handling – see multi-core features)
- Programmable Memory Management Unit (MMU)
 - 48-bit Virtual Addressing
 - 40-bit Physical Addressing – directly addresses up to 1 Terabyte
 - 1st level micro TLBs (uTLBs) – 16 entry instruction TLB, 32 entry data TLB
 - 2nd level TLBs – simultaneous access, variable and fixed page sizes
 - 64x2 entry VTLB, 512x2 entry 4-way set associative FTLB
 - Hardware table walk for fast page refills
- Power Management Features
 - Multi-core cluster power controller (CPC):
 - Register-based, visible to/controllable by operating system
 - Per CPU voltage domain gating; per CPU clock gating
 - Cluster level DVFS capable

- Core level
 - Course and fine-grained clock gating throughout core
 - Way prediction on data and instruction L1 caches
 - Instruction and register-based sleep modes
- EJTAG debug block and interface



Coherent Multi-Core Processor Features

- Superscalar, deeply OoO multi-core processor
- Complete multi-core system designed for maximum cluster-level bandwidth
 - Coherence manager- – supports multi-core configurations up to six cores in a single cluster
 - High-bandwidth 256-bit internal data paths and external system interface
 - Integrated L2 cache (L2\$): 4-way set associative, up to 8MB of memory
 - ECC option on L2\$ RAM for higher data reliability
 - Configurable wait states to RAM for optimal L2\$ design
 - L2\$ hardware pre-fetch for higher throughput and performance
- Up to two IO Coherence Units (IOCU) per coherent processing system
- Cluster Power Controller (CPC) for voltage/clock gating per-CPU
- 256-interrupt Global Interrupt Controller (GIC)
- Virtualization support at system level – IOcus have IO MMU, and GIC has virtualized interrupts

Specifications

Target	TSMC 28HPM
Frequency	1 GHz – 2+ GHz*
CoreMark/MHz (per core)	> 5
Total CoreMark @ 1.5GHz	> 7500 per core
DMIPS/MHz (per core)	3.5
Total DMIPS @ 1.5GHz	> 5250 per core

Notes: Frequencies indicated range from 12T SVt area-optimized in worst case silicon corner, to 12T MVt speed-optimized typical corner silicon. Final production RTL results may vary.

Each base core configuration:

- 32KB Data/Inst L1 caches with parity, BIST
- High-speed Integer + Floating Point (SP and DP) SIMD unit
- Fully-featured MMU, using multi-level TLB (I/D uTLBs + 128 entry VTLB + 1024 entry FTLB)

Multi-core cluster configuration:

- Dual fully-configured P6600 cores per above
- Coherence Manager + integrated 1MB L2\$ w/ECC
- One hardware IO Coherence Unit (IOCU) port

Implementation libraries/parameters – speed optimized, based on:

- TSMC 28HPM 12T standard cells + Synopsys memories
- Worst case, slow-slow corner silicon (zero temp, WCZ) with 8% OCV + 25ps clock jitter margins, except where noted at typical silicon

About Wave Computing

Wave Computing, Inc. is revolutionizing AI with its dataflow-based systems and solutions that deliver orders of magnitude performance improvements over legacy architectures. The company's vision is to bring deep learning to customers' data wherever it may be—from the edge to the datacenter—helping accelerate time-to-insight. Wave is powering the next generation of AI by combining its dataflow architecture with its MIPS embedded RISC multithreaded CPU cores and IP. Wave Computing has been named Frost & Sullivan's 2018 "Machine Learning Industry Technology Innovation Leader," and recognized by CIO Application Magazine's as one of the "Top 25 Artificial Intelligence Providers." Combined with MIPS, Wave now has over 400 granted and pending patents and hundreds of customers worldwide.

Wave Computing, the Wave Computing logo and MIPS are trademarks of Wave Computing, Inc. All other trademarks are used for identification purposes only and are the property of their respective owners.