

## Introduction

The PIC64GX microprocessor (MPU) is a 64-bit Linux® Operating System (OS)-capable processor that provides an innovative, mid-range, embedded compute platform that is based on the RISC-V® Instruction Set Architecture (ISA).

The PIC64GX MPU micro-architecture implementation is a simple, five-stage, single-issue, in-order pipeline that is not vulnerable to the Meltdown and Spectre exploits that frequently occur in common, out-of-order machines. The PIC64GX has five RISC-V cores which are coherent with the memory subsystem, allowing a versatile mix of deterministic real-time systems and Linux in a single, multi-core processor cluster. With integrated secure boot, Linux and Real-Time modes, a large and flexible L2 memory subsystem and a rich set of embedded peripherals, the PIC64GX MPU provides new choices for developing secure and power-efficient embedded compute platforms.

### Deterministic Asymmetric Multiprocessing (AMP) Linux Applications

Safety-critical, system control and security applications need the flexibility of the Linux OS and the determinism of real-time systems to control hardware. Typical Symmetric Multiprocessing (SMP) implementations may offer the flexibility of a rich operating system but are terrible for running real-time systems that need deterministic performance. PIC64GX MPUs feature a multi-core Linux-OS-capable processor that is coherent with the memory subsystem, allowing a versatile mix of deterministic real-time systems and the Linux OS in a single multi-core CPU cluster. PIC64GX MPUs enable you to create high-performance and hard real-time systems.

## Features

The PIC64GX1000 family of MPUs support the following features:

- Application Compute Complex With Real-Time Support

The PIC64GX1000 family of MPUs provides multi-core 64-bit RISC-V processing. Highlights of the Core Complex include:

- Four U54 64-bit RISC-V CPU cores
- Linux capable with AMP and deterministic latencies
- Simultaneous Linux and RTOS/bare metal in conjunction with E51 monitor processor
- Five-stage, single-issue, in-order pipeline
- Operates at up to 600 MHz
- Physical Memory Protection (PMP) unit
- Memory Management Unit (MMU)
- L1 memory subsystem with Single-Error Correct, Double-Error Detect (SECDED)
- 32 KB eight-way instruction cache or optional 28 KB tightly integrated memory
- 32 KB eight-way data cache
- 5K DMIPS performance

- Integrated Monitor Processor

The on-board monitor processor is an additional 64-bit RISC-V CPU core (E51) operating at 600 MHz

- 16 KB memory subsystem with SECDED configurable as two-way L1 instruction cache or as an instruction tightly integrated memory
- 8 KB data tightly integrated memory
- PMP unit

- Video Subsystem
  - The PIC64GX1000 MPU integrates MIPI CSI-2®, HDMI® 1.4 and a video pipeline.
    - Two lanes MIPI® CSI-2 at 1 Gbps for connection to cameras and sensors
    - Dedicated I<sup>2</sup>C, reset and standby signals for MIPI CSI-2 camera interfaces
    - HDMI 1.4 output with dedicated I<sup>2</sup>C and HPD signals for line drivers
    - Video pipeline with debayer, automatic white balance and color correction
- Memory Interfaces
  - DDR interface controller supporting DDR4-1333 or LPDDR4-1333
  - Up to 32 Gb per DDR interface
- On-chip Memory
  - 2 MB embedded SRAM for scratchpad memory
  - 128 KB of user Nonvolatile Memory (NVM) for boot Flash
  - 56 KB of secure NVM for user data and key storage
  - Coprocessor/Accelerator Interfaces
  - PCIe® Gen 2 x4 (FCV) or x1 (FCS) root ports
  - Peripheral Interfaces
  - Two SPI, five multi-mode UARTs, two I<sup>2</sup>C, 32 GPIOs, two CAN, MIPI CSI-2, HDMI 1.4, timers and watchdogs
- Defense-in-Depth Security
  - Dedicated secure enclave to support secure boot and platform root of trust
  - Full support for AES, SHA, HMAC, ECDSA, RSA, DSA and DRBG
  - Extensive anti-tamper detection and response capabilities

## Table of Contents

Introduction.....	1
Features.....	1
1. Overview.....	5
2. Product Family Overview.....	6
3. Processor System.....	7
3.1. E51 RISC-V Monitor Core.....	7
3.2. U54 RISC-V Application Cores.....	9
4. Debug.....	12
4.1. CPU Debug.....	12
5. Interrupts.....	13
6. Memory Subsystems.....	14
6.1. L2 Memory Subsystem.....	14
6.2. DDR Memory Controller.....	14
6.3. Processor Interconnect.....	15
7. Peripherals.....	16
7.1. Gigabit Ethernet MAC.....	16
7.2. PHY Interfaces.....	16
7.3. MMC 5.1/SD/SDIO/eMMC.....	17
7.4. USB 2.0 OTG.....	18
7.5. User Crypto.....	18
7.6. Controller Area Network.....	19
7.7. QSPI XIP Controller.....	19
7.8. Serial Peripheral Interface.....	20
7.9. Multi-Mode UART.....	20
7.10. Real Time Counter.....	21
7.11. Watchdog Timer.....	21
7.12. Timer.....	21
7.13. Peripheral Memory SECDED Reporting and Error Injection.....	21
7.14. DMA Controller.....	22
8. Boot Process.....	23
9. PCI Express.....	24
9.1. PCI Express Features.....	24
10. Tools.....	25
11. Compatibility with PolarFire SoC.....	26
12. Revision History.....	27
Microchip Information.....	28
The Microchip Website.....	28

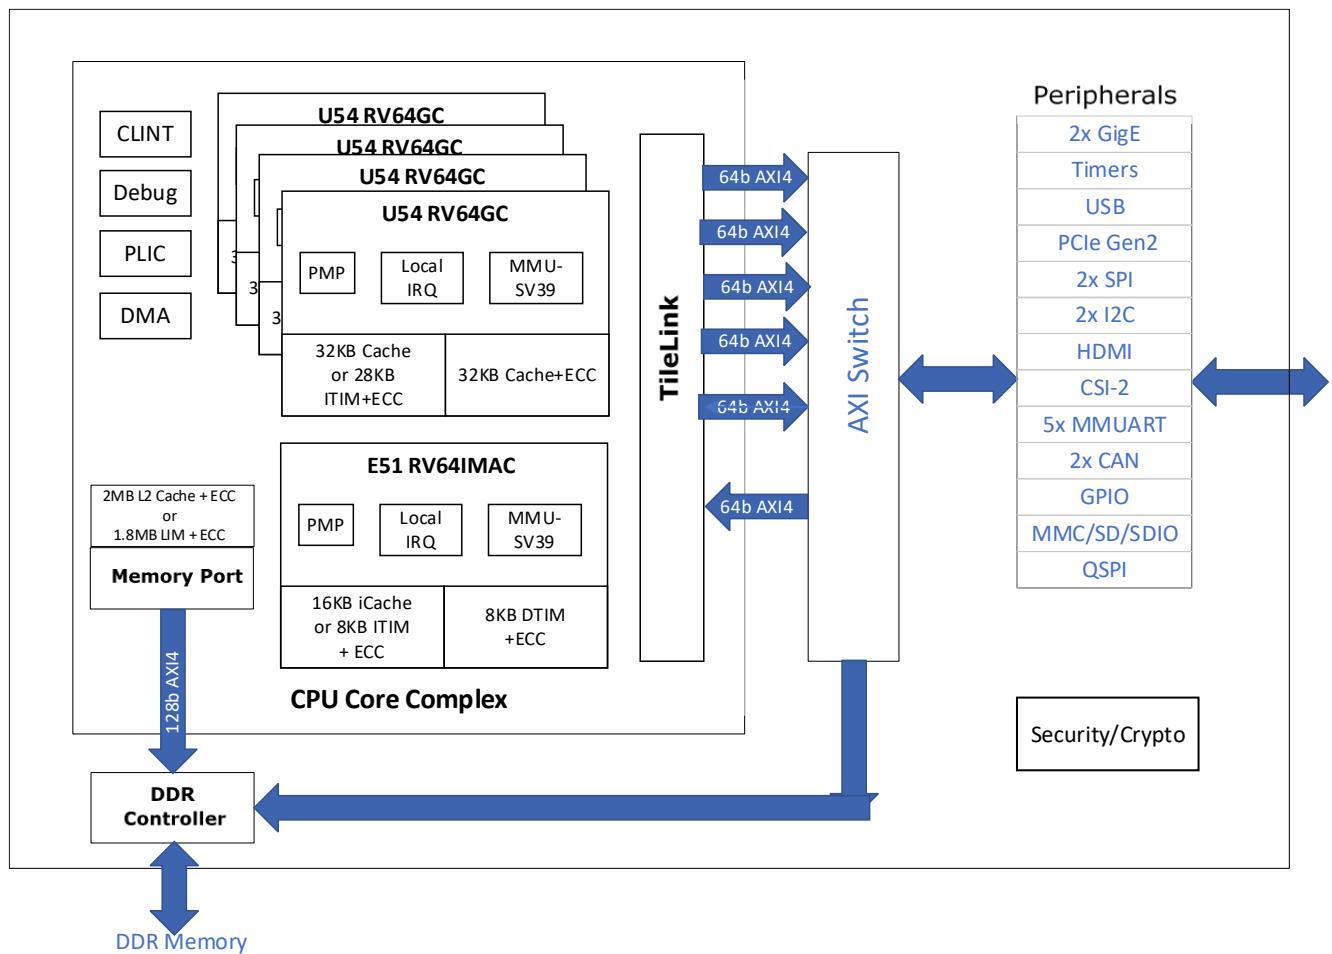
Product Change Notification Service.....	28
Customer Support.....	28
Microchip Devices Code Protection Feature.....	28
Legal Notice.....	28
Trademarks.....	29
Quality Management System.....	30
Worldwide Sales and Service.....	31

## 1. Overview

PIC64GX1000 is a highly power-efficient 64-bit Linux capable processor, innovative, mid-range, embedded compute platform based on the RISC-V ISA.

The RISC-V CPU micro-architecture implementation is a simple, five-stage single issue, in-order pipeline that does not suffer from the Meltdown and Spectre exploits found in common out-of-order machines. It has five RISC-V cores which are coherent with the memory subsystem allowing a versatile mix of deterministic real time systems and Linux in a single, multi-core processor cluster. With Secure Boot built-in, innovative Linux and Real Time modes, a large Flexible L2 memory subsystem, and a rich set of embedded peripherals. The PIC64GX1000 MPU provides embedded developers new choices in secure, power-efficient, embedded compute platforms.

**Figure 1-1. PIC64GX1000 Detailed Block Diagram**



## 2. Product Family Overview

The PIC64GX1000 comes in two package-based derivatives with specified pinouts, I/O counts, and features. The devices are available in several temperature ranges based on the ordering part number.

**Table 2-1.** Product Family Table

Features	PIC64GX1000-V/FCS	PIC64GX1000-V/FCV
U54 RISC-V Application Cores	4	4
U51 RISC-V Monitor Core	1	1
PCIe Gen2 host	X1	X4
HDMI	1	1
MIPI CSI-2	X2	X2
DDR Data Bus	16	32
Ethernet MAC	2	2
DPA Resistant Data Security	Y	Y
USB 2.0 OTG	1	1
SD/MMC	1	1
CAN 2.0 A/B	<b>2</b>	2
QSPI	<b>4ch</b>	4ch
SPI	<b>2</b>	<b>2</b>
I2C	<b>2</b>	2
MMUART	<b>5</b>	5
Type/Size/Pitch (Commercial/Industrial)	Total I/O	
FCSG325 (11 mm x11 mm , 0.5 mm)	200	N/A
FCVG484 (19 mm x 19 mm, 0.8 mm)	N/A	244

### 3. Processor System

This section discusses the processor system.

#### 3.1 E51 RISC-V Monitor Core

The following table describes the features of E51.

**Table 3-1.** E51 RISC-V Monitor Core Features

Feature	Description
ISA	RV64IMAC
iCache/ITIM	16 KB 2-way set-associative/8 KB ITIM
DTIM	8 KB
ECC Support	Single-Error Correction and Double-Error Detection (SECDED) on iCache and DTIM.
Modes	Machine Mode, User Mode

Typically, in a system, the E51 is used to execute the following functions:

- Bootloader to boot the operating system on U54 cores
- Bare-metal user applications
- Monitoring user applications on U54 cores

**Note:** Load-Reserved and Store-Conditional atomic instructions (lr, sc) are not supported on the E51 processor core.

##### 3.1.1 E51 Instruction Fetch Unit

The E51 instruction fetch unit consists of a two-way, 16 KB instruction cache that supports 64-byte cache lines. The access latency is one clock cycle. Writes to memory may be synchronized with the instruction fetch stream with a `FENCE.I` instruction. The branch predictor comprises a branch target buffer (BTB), which predicts the target of taken branches and jumps; a branch history table (BHT), which predicts the direction of conditional branches; and a return-address stack (RAS), which predicts the target of procedure returns. The BTB is configured to hold 40 entries. The RAS is configured to hold two entries. The BHT uses a gshare prediction scheme with 7 bits of global history to access an array of 128, two-bit saturating counters. The branch predictor has a one cycle latency, so that correctly predicted control-flow instructions result in no penalty. The branch predictor can be turned off during device configuration to create deterministic systems.

##### 3.1.2 E51 I-Cache Reconfiguration

The instruction cache can be partially reconfigured into an Instruction Tightly Integrated Memory (ITIM), which occupies a fixed address range in the memory map. ITIM provides high performance, predictable instruction delivery. Fetching an instruction from ITIM is as fast as an instruction-cache hit, with no possibility of a cache miss. ITIM can hold data as well as instructions, though loads and stores to ITIM are not as performant as loads and stores to DTIM. The instruction cache can be configured as ITIM for all ways except for 1 in units of cache lines (64 bytes). A single cache way must remain as an instruction cache. ITIM is allocated simply by storing to it. A store to the  $n^{\text{th}}$  byte of the ITIM memory map reallocates the first  $n+1$  bytes of instruction cache as ITIM, rounded up to the next cache line. ITIM is deallocated by storing zero to the first byte after the ITIM region. The deallocated ITIM space is automatically returned to the instruction cache. For determinism, software must clear the contents of ITIM after allocating it. It is unpredictable whether ITIM contents are preserved between deallocation and allocation. ITIM is typically used for reduced latency requirements like for an ISR.

##### 3.1.3 E51 Execution Pipeline

The E51 execution unit is a single-issue, in-order pipeline. The pipeline comprises five stages:

1. Instruction fetch
2. Instruction decode and register fetch
3. Execute
4. Data memory access
5. Register write back

The pipeline has a peak execution rate of one instruction per clock cycle. It is fully bypassed, so that most instructions have an apparent one-cycle result latency. There are several exceptions:

LD and LW have a two-cycle result latency, assuming a cache hit.

LH, LHU, LB, and LBU have a three-cycle result latency, assuming a cache hit.

MUL, MULW, MULH, MULHU, MULHSU, DIV, DIVU, REM, REMU, DIVW, DIVUW, REMW, and REMUW have between a 2-cycle and 66-cycle result latency, depending on operand values.

CSR reads have a three-cycle result latency.

The pipeline only interlocks on read-after-write and write-after-write hazards, so instructions may be scheduled to avoid stalls.

The iterative multiplier is configured to produce 16 bits per cycle with an early-out option. The iterative divider has latency of between three and 66 cycles and an early-out option.

Branch and jump instructions transfer control from the memory access pipeline stage. Correctly predicted branches and jumps incur no penalty, whereas unpredicted branches and jumps incur a three-cycle penalty. Most CSR writes result in a pipeline flush, a five-cycle penalty.

### **3.1.4 E51 Data Memory System**

The E51 data memory system consists of 8 KB Data Tightly Integrated Memory (DTIM). The access latency is two clock cycles for full words and three clock cycles for smaller quantities. Misaligned accesses are not supported in hardware and result in a trap to support software emulation. Stores are pipelined and commit on cycles where the data memory system is otherwise idle. Loads to addresses currently in the store pipeline result in a five-cycle penalty.

### **3.1.5 E51 Memory Single-Bit Errors**

When a single bit error is detected in the E51 L1 memory used as a cache, the error is corrected and the cache line is flushed consequently writing to the next level in the memory hierarchy (L2 cache). When a single bit error is detected in the E51 L1 memory configured as an ITIM, the error is corrected and written back to the ITIM location.

### **3.1.6 E51 Memory Error Correction**

The E51 DTIM implements single-error correcting, double-error detecting (SECDED) error correcting code (ECC). The granularity at which this protection is applied (the codeword) is 32 bits (with an ECC overhead of 7 bits per codeword).

### **3.1.7 E51 Memory Error Reporting**

ECC events are reported by the Bus-Error Unit (BEU) for a given core. The BEU can be configured to generate interrupts either globally to the Platform Level Interrupt Controller (PLIC), or locally to the specific part where the ECC event occurred. When BEU interrupts are enabled, software can then be used to monitor and count ECC events. In order to detect uncorrectable ECC errors in the L1 memory system, interrupts must be enabled in the BEU. Specifically, to halt execution of a core when an uncorrectable instruction is detected, the BEU must be configured to generate a local interrupt.

### 3.1.8 E51 Local Interrupts

The E51 supports up to 48 local interrupt sources routed directly to the core. The E51 core receives the same 48 interrupts as the U54 cores.

## 3.2 U54 RISC-V Application Cores

The following table describes the features of the U54 application cores.

**Table 3-2. U54 RISC-V Application Cores Features**

Feature	Description
ISA	RV64IMAFDC (RV64GC)
Instruction cache	32 KB, 8-way
Instruction tightly integrated memory (ITIM)	Maximum of 28 KB
Data cache	32 KB, 8-way
ECC support	Single-error correct, double-error detect on the instruction cache/ITIM and data cache
Virtual memory support	The U54 has support for Sv39 virtual memory support with a 39-bit virtual address space, 38-bit physical address space, and 32-entry TLB
Privileged modes	Machine (M), Supervisor (S), User (U)

### 3.2.1 U54 Instruction Memory System

The instruction memory system consists of a dedicated 32 KB, 8-way, set-associative, Virtually Indexed Physically Tagged (VIPT) instruction cache. The access latency of all blocks in the instruction memory system is one clock cycle. The instruction cache is not kept coherent with the rest of the platform memory system. Writes to instruction memory must be synchronized with the instruction fetch stream by executing a FENCE.I instruction. The instruction cache has a line size of 64 bytes and a cache line fill will trigger a burst access outside of the PIC64GX1000 CPU core complex. The core will cache instructions from executable addresses, with the exception of the ITIM. Trying to execute an instruction from a non-executable address will result in a synchronous trap.

### 3.2.2 U54 I-Cache Reconfiguration

The instruction cache can be partially reconfigured into an Instruction Tightly Integrated Memory (ITIM), which occupies a fixed address range in the memory map. ITIM provides high-performance, predictable instruction delivery. Fetching an instruction from ITIM is as fast as an instruction-cache hit, with no possibility of a cache miss. ITIM can hold data as well as instructions, though loads and stores to ITIM are not as performant as loads and stores to DTIM. The instruction cache can be configured as ITIM for all ways except for 1 in units of cache lines (64 bytes). A single-instruction cache way must remain as an instruction cache. ITIM is allocated simply by storing to it. A store to the  $n^{\text{th}}$  byte of the ITIM memory map reallocates the first  $n+1$  bytes of instruction cache as ITIM, rounded up to the next cache line. ITIM is deallocated by storing zero to the first byte after the ITIM region. The deallocated ITIM space is automatically returned to the instruction cache. For determinism, software must clear the contents of ITIM after allocating it. It is unpredictable whether ITIM contents are preserved between deallocation and allocation.

### 3.2.3 U54 Instruction Fetch Unit

The U54 instruction fetch unit contains branch prediction hardware to improve performance of the processor core. The branch predictor comprises a 40-entry branch target buffer (BTB) that predicts the target of taken branches, a 128-entry branch history table (BHT) that predicts the direction of conditional branches, and a 2-entry return-address stack (RAS) that predicts the target of procedure returns. The branch predictor has a one-cycle latency, so that correctly predicted control-flow instructions result in no penalty. Unpredicted control-flow instructions incur a three-cycle penalty. The branch predictor can be turned off during device configuration to create deterministic systems. The U54 implements the standard Compressed (C) extension to the RISC-V architecture which allows for 16-bit RISC-V instructions.

### 3.2.4 U54 Execution Pipeline

The U54 execution unit is a single-issue, in-order pipeline. The pipeline comprises five stages: instruction fetch, instruction decode and register fetch, execute, data memory access, and register writeback.

The pipeline has a peak execution rate of one instruction per clock cycle, and is fully bypassed so that most instructions have a one-cycle result latency. There are several exceptions:

- LW has a two-cycle result latency, assuming a cache hit.
- LH, LHU, LB, and LBU have a three-cycle result latency, assuming a cache hit.
- CSR reads have a three-cycle result latency.
- MUL, MULH, MULHU, and MULHSU have a 5-cycle result latency.
- DIV, DIVU, REM, and REMU have between a 2-cycle and 33-cycle result latency, depending on the operand values.

The pipeline only interlocks on read-after-write and write-after-write hazards, so instructions may be scheduled to avoid stalls.

The U54 implements the standard Multiply (M) extension to the RISC-V architecture for integer multiplication and division. The U54 has a 16-bit per cycle hardware multiply and a 4-bit per cycle hardware divide.

Branch and jump instructions transfer control from the memory access pipeline stage. Correctly predicted branches and jumps incur no penalty, whereas unpredicted branches and jumps incur a three-cycle penalty.

Most CSR writes result in a pipeline flush with a five-cycle penalty.

### 3.2.5 U54 Data Memory System

The U54 data memory system has a 8-way set-associative 32 KB write-back data cache with 64 B cache lines. The data cache is Virtually Indexed Physically Tagged (VIPT). Access latency is two clock cycles for words and double-words, and three clock cycles for smaller quantities. Misaligned accesses are not supported in hardware and result in a trap to support software emulation. The data caches are kept coherent with a directory-based cache coherence manager, which resides in the outer L2 cache. Stores are pipelined and commit on cycles where the data memory system is otherwise idle. Loads to addresses currently in the store pipeline result in a five-cycle penalty

### 3.2.6 U54 Atomic Operations

The U54 core supports the RISC-V standard Atomic (A) extension on regions of the memory map denoted by the attribute A. Atomic memory operations to regions that do not support them generate an access exception precisely at the core. The load-reserved and store-conditional instructions are only supported on cached regions; therefore, generate an access exception on DTIM and other uncached memory regions. See *The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Version 2.1 [1]* for more information on the instructions added by this extension.

### 3.2.7 U54 Floating Point Unit (FPU)

The U54 FPU provides full hardware support for the IEEE® 754-2008 floating-point standard for 32-bit, single-precision and 64-bit, double-precision arithmetic. The FPU includes a fully pipelined fused-multiply-add unit and an iterative divide and square-root unit, magnitude comparators, and float-to-integer conversion units, all with full hardware support for subnormals and all IEEE default values.

### 3.2.8 U54 Virtual Memory Support

The U54 has support for virtual memory using a Memory Management Unit (MMU). The MMU supports the Bare and Sv39 modes as described in the *RISC-V Instruction Set Manual, Volume II*:

*Privileged Architecture, Version 1.10* [2]. The U54 MMU has a 39-bit virtual address space mapped to a 38-bit physical address space. A hardware page-table walker refills the address translation caches. Both instruction and data address translation caches are fully associative, and have 32 entries. The MMU supports 2 MB megapages and 1 GB gigapages to reduce translation overheads for large contiguous regions of virtual and physical address space. Note that the U54 does not automatically set the Accessed (A) and Dirty (D) bits in a Sv39 Page Table Entry (PTE). Instead, the U54 MMU will raise a page fault exception for a read to a page with PTE.A=0 or a write to a page with PTE.D=0.

### 3.2.9 U54 Local Interrupts

Each U54 supports up to 48 local interrupt sources that are routed directly to the core. These are shared with the U51 core

### 3.2.10 U54 Memory Error Correction

The U54 ITIM and DTIM implement single-error correcting, double-error detecting (SECDEC) error correcting code (ECC). The granularity at which this protection is applied (the codeword) is 32-bit (with an ECC overhead of 7 bits per codeword).

#### 3.2.10.1 U54 Memory Single-Bit Errors

When a single bit error is detected in the U54 L1 memory used as a cache, the error is corrected and the cache line is flushed consequently writing to the next level in the memory hierarchy (L2 cache). When a single bit error is detected in the U54 L1 memory configured as an ITIM, the error is corrected and written back to the ITIM location.

#### 3.2.10.2 U54 Memory Error Reporting

ECC events are reported by the Bus-Error Unit (BEU) for a given core. The BEU can be configured to generate interrupts either globally to the Platform Level Interrupt Controller (PLIC), or locally to the specific part where the ECC event occurred. When BEU interrupts are enabled, software can then be used to monitor and count ECC events. In order to detect uncorrectable ECC errors in the L1 memory system, interrupts must be enabled in the BEU. Specifically, to halt execution of a core when an uncorrectable instruction is detected, the BEU must be configured to generate a local interrupt. Uncorrectable ECC errors in the L1 system are reported as a HALT\_CPU\_n signal, where n = 0 – 5 and n = 1 through 5 = the U54 application cores.

### 3.2.11 Physical Memory Protection

Each CPU in the PIC64GX1000 includes a physical memory protection (PMP) unit compliant with the *RISC-V Instruction Set Manual, Volume II: Privileged Architecture, Version 1.10*. The PMP unit can be used to set memory access privileges (read, write, execute) for specified memory regions. Each PMP supports 16 regions with a minimum region size of 4 bytes.

#### 3.2.11.1 Functional Description

PIC64GX1000 includes a Physical Memory Protection (PMP) unit, which can be used to restrict access to memory and isolate processes from each other. The PMP unit has 16 regions and a minimum granularity of 4 bytes. It is permitted to have overlapping regions. The PIC64GX1000 PMP unit implements the architecturally defined pmpcfgX CSRs pmpcfg0 and pmpcfg2 supporting 16 regions. pmpcfg1 and pmpcfg3 are implemented but hardwired to zero.

#### 3.2.11.2 Region Locking

The PMP allows for region locking whereby once a region is locked, further writes to the configuration and address registers are ignored. Locked PMP entries may only be unlocked with a system reset. A region may be locked by setting the L bit in the pmpcfg register. In addition to locking the PMP entry, the L bit indicates whether the R/W/X permissions are enforced on M-Mode accesses. When the L bit is set, these permissions are enforced for all privilege modes. When L bit is clear, the R/W/X restrictions apply only to U-mode.

## 4. Debug

This section discusses the debug process.

### 4.1 CPU Debug

Each CPU has up to ten breakpoint registers. Breakpoints can halt the respective CPU under the following conditions:

- Address match on Load
- Address match on Store
- Address match on Instruction Fetch
- Address match on User mode
- Address match on Supervisor mode
- Address match on Machine mode

A successful match on address can generate an exception or place the machine into debug mode.

#### 4.1.1 Trace

Trace includes an instruction trace interface module. For each core, the following properties can be captured when an instruction is retired or trapped and trace is enabled:

- The address of the instruction
- The instruction
- Privileged mode during execution
- Trap or retired indication
- Interrupt or exception indication
- Exception cause
- Exception data

#### 4.1.2 AXI Bus Monitors

There are two AXI bus monitors within the Core Complex that allow, at run time, observation of AXI transactions. They can be used for debugging, reporting diagnostics, performance profiling, bare metal security, and other applications. For example, the AXI monitors can passively filter on specific addresses, which master is performing the R/W, without affecting traffic at run time.

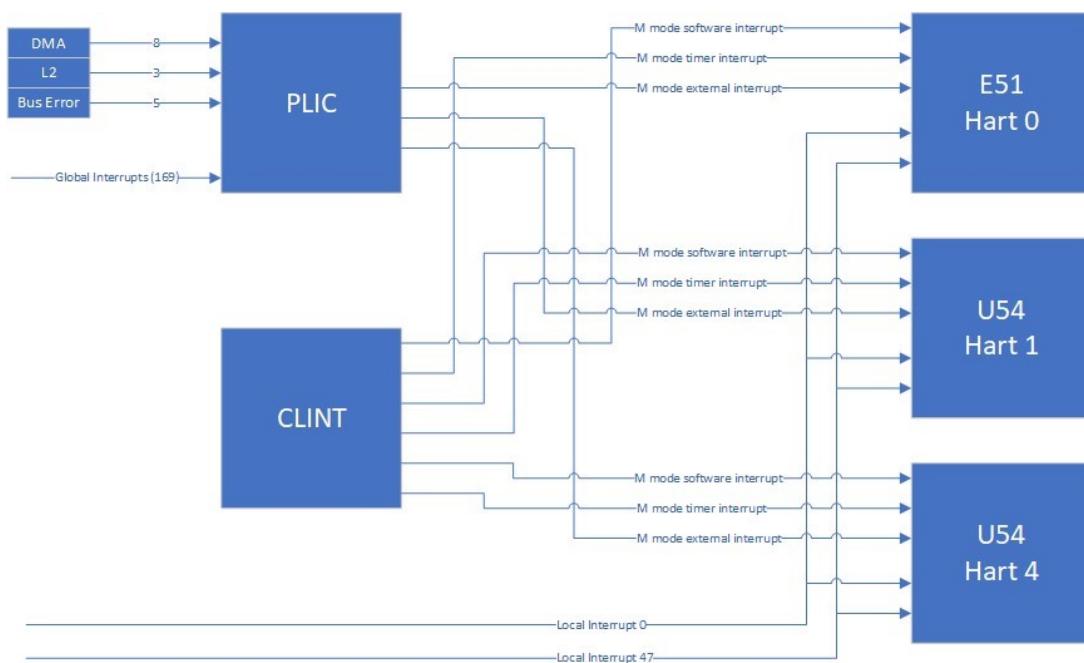
The AXI-64 bus monitor is configured to provide full address and data trace on the slave side of the AXI switch.

The AXI-128 bus monitor is configured to provide full address trace on the AXI4-128 bus connecting the Core Complex L2 interface to the DDR memory. Data trace is not supported in this case. This allows the ability to trace the effectiveness of the cache and DDR response rates

## 5. Interrupts

The PIC64GX1000 has support for the following interrupts: local (including software and timer) and global. Local interrupts are signaled directly to an individual hart with a dedicated interrupt value. This allows for reduced interrupt latency as there is no arbitration required to determine which hart will service a given request, nor additional memory accesses required to determine the cause of the interrupt. Software and timer interrupts are local interrupts generated by the Core Local Interruptor (CLINT). Global interrupts by contrast, are routed through a Platform-Level Interrupt Controller (PLIC), which can direct interrupts to any hart in the system via the external interrupt. Decoupling global interrupts from the hart(s) allows the design of the PLIC to be tailored to the platform, permitting a broad range of attributes like the number of interrupts and the prioritization and routing schemes. By default, all interrupts are handled in machine mode. The U54s, which support supervisor mode, can selectively delegate interrupts to supervisor mode.

**Figure 5-1.** Interrupt System



## 6. Memory Subsystems

PIC64GX1000 contains an on-chip 128 KB embedded non-volatile memory (eNVM) for user code, a flexible L2 memory subsystem, and an integrated DDR memory controller.

### 6.1 L2 Memory Subsystem

The PIC64GX1000 Level 2 Cache Controller is used to provide access to fast copies of memory for masters in a core complex. The Level 2 Cache Controller also acts as a directory-based coherency manager. The Level 2 Cache Controller offers extensive flexibility as it allows for several features in addition to the Level 2 Cache functionality, such as memory-mapped access to L2 Cache RAM for disabled cache ways, scratchpad functionality, way masking and locking, and ECC support with error tracking statistics, error injection, and interrupt signaling capabilities.

The L2 Cache Controller consists of 4 banks where each bank contains 512 sets of 16 ways, and each way contains a 64-byte block. This subdivision into banks helps facilitate increased available bandwidth between CPU masters and the L2 Cache, as each bank has its own 128-bit inner port. As such, multiple requests to different banks may proceed in parallel. The outer port of the L2 Cache Controller is a 128-bit port shared among all banks and is connected to the DDR controller.

When cache ways are disabled, they are addressable in the L2 Loosely Integrated Memory (L2LIM) address space as described in the [PIC64GX1000 Data Sheet](#). Fetching instructions or data from the L2-LIM provides deterministic behavior equivalent to an L2 cache hit, with no possibility of a cache miss. Accesses to L2-LIM are always given priority over cache way accesses which target the same L2 cache bank. Out of reset all ways, except for way 0, are disabled. Cache ways can be enabled by writing to specific control registers. Once a cache way is enabled, it can not be disabled unless the CPU complex is reset. The highest numbered L2 cache way is mapped to the lowest L2-LIM address space, and way 1 occupying the highest L2-LIM address range. As L2 cache ways are enabled, the size of the L2-LIM address space shrinks.

The L2 Cache Controller has a dedicated scratchpad address region that allows for allocation into the cache using an address range which is not memory backed. This address region is denoted as the L2 Zero Device in the memory map. Writes to the scratchpad region will allocate into cache ways that are enabled and not masked. Care must be taken with the scratchpad, however, as there is no memory backing this address space. Cache evictions from addresses in the scratchpad will result in data loss. The main advantage of the L2 Scratchpad over the L2-LIM is that it is a cacheable region allowing for data stored to the scratchpad to also be cached in a master's L1 data cache resulting in faster access.

### 6.2 DDR Memory Controller

The hardened, 16-bit DDR memory controller supports the following features:

- DDR4, LPDDR4, 16-bit, or 32-bit memory support
- Dual Rank support for dual die packages
- Max rate support 1333 Mbps
- DDR memory test feature
- Reorder queue to optimize DDR performance
- Two AXI interfaces
  - 128-bit from CPU L2 Cache
  - 64-bit from central AXI switch
- 32 outstanding transactions per AXI interface
- An integrated Clock Domain Crossing (CDC) circuit allowing the DDR controller clock to be independent of the CPU clock

- A dedicated PLL for DDR clock generation

It can support up to 8 GB of external DDR4 memory. SECDED capability is also provided when configured to a 36-bit bus width, as listed in the following table.

**Table 6-1.** DDR Memory Controller

Configuration	Active Pads	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4 <sup>1</sup>
5x8 DDR with SECDED	36	DDRx8	DDRx8	DDRx8	DDRx8	DDRx8 (4 used)
4x8 DDR	32	DDRx8	DDRx8	DDRx8	DDRx8	Not used
3x16 DDR with SECDED	36	DDRx16		DDRx16		DDRx16 (4 used)
2x16 DDR	32	DDRx16		DDRx16		Not used
3x16 DDR with SECDED	18	DDRx8	DDRx8	—	—	DDRx8 (2 used)
2x16 DDR	16	DDRx8	DDRx8	—	—	—
1x16 DDR with SECDED	18	DDRx16		—	—	DDRx16 (2 used)
1x16 DDR	16	DDRx16		—	—	—

**Note:** 1. Lane 4 is only 4 bits wide, the upper data bits on the DDR memory are not connected.

## 6.3

### Processor Interconnect

There are two interconnect switches built into the PIC64GX1000. First, there is a fully populated coherent switch that manages coherence through the memory subsystems and provides a deterministic data path to the L2 memory subsystem when it is configured as a loosely integrated memory. Additionally, a central AMBA I/O switch manages the interconnect between the CPU complex, the peripheral I/O space and the hardened DDR memory controller. The AMBA switch also includes Quality of Service (QoS) features. The QoS feature is essentially a 4-bit value denoting priority for the data path. The central I/O switch is partially connected and supports 15 masters with nine slaves. The AMBA switch also contains a Memory Protection Unit (MPU) scheme that mimics the Physical Memory Protection (PMP) scheme defined in the RISC-V Privileged Specification.

Specifically, the bus masters listed in the following table pass through the AMBA MPU. PMP region address granularity start at 4096 and increase by powers of two. Regions can be defined to support execution, read or write. An additional lock bit can be set on each register that prevents further modification to the MPU protection scheme until the next power on reset.

## 7. Peripherals

The PIC64GX1000 supports many peripherals apart from DDR IOs, SGMII IO for Ethernet MACs, and IOs for reference clocks:

- eMMC/SD/SDIO
- USB
- QSPI-XIP
- Two CAN
- Five UARTs
- Two SPI
- Two I2C
- GPIO
- MIPI CSI-2 (2 lanes)
- PCIe Gen2
- 1G Ethernet x2

All these IOs are bonded out to pins in all PIC64GX1000 packages.

### 7.1 Gigabit Ethernet MAC

PIC64GX1000 contains two identical Gigabit Ethernet MACs (GEM) integrated into the Core Complex. Each MAC can contain a maximum frame length of 10,240 bytes that are SECDED protected. Additionally, the GEM supports a built-in packet buffer DMA. Each GEM supports the following IEEE® 802 standards:

- IEEE 802.3br Frame Pre-Emption (or Interspersing Express Traffic)
- IEEE 802.1Qci Receive (Ingress) Traffic Policing
- IEEE 802.1Qbb Priority-Based Flow Control
- IEEE 802.1Q VLAN Tagging with Recognition of Incoming VLAN and Priority Tagged Frames
- IEEE 802.1AS
- IEEE 802.1Qav
- IEEE 802.1Qbv
- IEEE 1588-2002 (v1), IEEE 1588-2008 (v1 and v2)
- IEEE 802.1CB Frame Redundancy and Elimination

### 7.2 PHY Interfaces

Each GEM is configured to simultaneously support TBI and GMII modes. When using TBI, the PCS block of the MAC is used, but not as an IEEE802.3X interface to a transceiver, but rather, is fed into a dedicated SERDES block and from there, interfaces to a PHY. This serialized interface between MAC and PHY is known as SGMII. In SGMII mode, the PCS interface and the link speed auto-negotiation blocks are re-purposed from their 802.3X function and are instead used to convey control information related to the MAC-PHY interface.

#### 7.2.1 Direct SGMII I/O

The following functionality provides an SGMII interface from the GEM to the built-in SGMII PHY:

- Clock Domain Recovery (CDR) of received 125 MHz clock
- Serializing/De-serializing

- PLL for synthesis of a 125 MHz transmit clock
- I/O buffers (four I/Os per MAC instance) allowing for differential transmit and receive data pairs

**Note:** It is not possible to support the SyncE protocol when using the direct SGMII interface.

### 7.2.2 PHY Management

Each MAC has an MDC output and an MDIO input and output port, which is brought out separately for each MAC instance. If desired, however, the user could utilize only one management interface (and not use the second), as it is possible to control multiple PHYs using the one interface (if hardware separation is not required).

### 7.2.3 Internal Filtering

The GEM is configured to have four internal specific address filters configured. Each filter can be configured to contain a MAC address, which can be specified to be compared against the source address (SA) or destination address (DA) of each received frame. There is also a mask field to allow certain bytes of the address to be excluded in the comparison. If the filtering matches for a specific frame, then it is passed on to the DMA memory. Otherwise the frame is dropped. Frames may also be filtered using the Type ID field for matching. There are four Type ID registers in the internal register space and these may be enabled individually. Hashing of the received frame's DA may be configured, as described in the PIC64GX1000 Datasheet

### 7.2.4 External Filtering

To allow for more sophisticated matching of incoming frames, based on more than just addresses, for example, an external filtering interface is present. As a frame is received, the MAC parses the frame and determines what field is currently present. A strobe signal is provided to allow latching of each field in the fabric. Customized (combinational) matching circuitry can then analyze whether or not it is interested in the particular frame (filtering use case) or what receive priority queue to put the frame in (prioritizing use case).

## 7.3 MMC 5.1/SD/SDIO/eMMC

PIC64GX1000 contains one MMC5.1 compliant peripheral and PHY. The I/Os do not support the dynamic voltage scaling of some SD cards, external voltage level translators should be used if required. The following eMMC/SD card standards are supported in PIC64GX1000. The SDIO/eMMC interface supports its own DMA controllers for data transfers. The DMA controllers support SDMA and ADMA2 modes.

### 7.3.1 SD Card Standards

The following is a list of SD card standards:

- Default Speed (DS)
- High Speed (HS)
- UHS-I SDR12
- UHS-I SDR25
- UHS-I SDR50
- UHS-I SDR104
- UHS-I DDR50

### 7.3.2 eMMC Standards

The following is a list of eMMC card standards:

- Standard Speed
- High Speed

- DDR52
- HS200
- HS400
- HS400 Enhanced Strobe

## 7.4 USB 2.0 OTG

PIC64GX1000 includes a USB 2.0 OTG-compliant core with an ULPI interface. The USB core supports the following features:

- Operates either as the function controller of a high/full-speed USB peripheral or as the host/peripheral in point-to-point or multi-point communications with other USB functions
- Complies with the USB 2.0 standard for high-speed (480 Mbps) functions and with the on-the-go supplement to the USB 2.0 specification
- Supports OTG communications with one or more high-, full-, or low-speed device
- Supports Session Request Protocol (SRP) and Host Negotiation Protocol (HNP)
- Supports suspend and resume signaling
- Supports Link Power Management
- Offers dynamic allocation of endpoints to maximize number of devices supported
- Number of Tx endpoints: 4 (plus control endpoint)
- Number of Rx endpoints: 4 (plus control endpoint)
- Multipoint capabilities supported
- Software connect/disconnect feature enabled
- High bandwidth support for ISO enabled endpoints
- Hardware-selectable option for 8-bit/4-bit LPI interface
- Vendor control and status register enabled with width of 4 and 8, respectively
- Number of DMA channels: 4
- Dynamic FIFO support allows dynamic allocation of buffer depth per enabled endpoint

## 7.5 User Crypto

PIC64GX1000 embeds an Athena TeraFire F5200B side channel resistant crypto coprocessor. Internal memories in the Athena core are SECDED protected.

### 7.5.1 TeraFire EXP-F5200B Supported Protocols Features

The following is a list of TeraFire EXP-F5200B supported protocols features:

- TRNG SP800-90A CTR\_DRBG-2565; SP800-90B (draft) NRBG
- AES 128/192/256 key lengths E/D (ECB, CBC, CTR, OFB, CFB, GCM, KeyWrap)
- SHA-1/224/256/384/512
- HMAC-SHA-1/224/256/384/512; GMAC-AES; CMAC-AES
- SHA-256 Key Tree
- ECC: NIST P256/384/521 and Brainpool P256/384/512 curves; KeyGen, KAS - ECC CDH, ECDSA SigGen & SigVer, PKG, PKV
- IFC: 1024/1536/2048/3072/4096 RSA E/D; SSA\_PKCS1\_V1\_5 SigGen & SigVer; ANSI X9.31 SigGen & SigVer
- FFC: 1024/153/2048/3072/4096; KAS - DH, DSA SigGen & SigVer

## 7.6 Controller Area Network

PIC64GX1000 integrates two controller area network cores compliant to CAN 2.0 A and B and conforms to ISO 11898-1. Internal message SRAM is SECDED-protected. The following are CAN controller features.

### 7.6.1 Receive Path

The following is a list of properties of receive path:

- 32 receive buffers
- Each buffer has its own message filter
- Message filter covers: ID, IDE, RTR, data byte 1, and data byte 2
- Message buffers can be linked together to build a bigger message array
- Automatic remote transmission request (RTR) response handler with optional generation of RTR interrupt

### 7.6.2 Transmit Path

The following is a list of properties of transmit path:

- 32 Tx message holding registers with programmable priority arbitration
- Message abort command
- Single-shot transmission (no automatic re-transmission upon error or arbitration loss)

### 7.6.3 Debugging Support

The following is a list of features of debugging support:

- Listen Only mode
- Internal Loop Back mode
- External Loop Back mode
- Error capture register
- SRAM test mode to support software based memory testing (SRAM is addressable by the CPUs if the core is disabled)

## 7.7 QSPI XIP Controller

PIC64GX1000 integrates a Quad SPI (QSPI) Flash controller with eXecute In Place (XIP) capabilities. The following is a list of QSPI features:

- Master SPI Data Rate
  - Programmable SPI clock HCLK/2, HCLK/4 or HCLK/6
  - Maximum data rate is HCLK/2
- FIFOs
  - Transmit and Receive FIFO
  - 16-byte transmit FIFO depth
  - 32-byte receive FIFO depth
- SPI Protocol
  - Master operation
  - Motorola SPI supported
  - Slave select operation in idle cycles configurable
  - Supports extended SPI operation (1, 2, and 4 bits)

- Supports QSPI operation (4-bit operation)
- Supports BSPI operation (2-bit operation)
- Support XIP (execute in place)
- Supports 3 or 4-byte SPI address.
- Frame Size
  - Supports 8-bit frames directly
  - Back-to-back frame operation supports >8-bit frames
  - Supports up to 4GB Transfer ( $2^{**32}$  bytes)
- Processor Overhead Reduction
  - Support of SPI Flash command/data packets with automatic data generation and discard function

## 7.8 Serial Peripheral Interface

Serial peripheral interface (SPI) is a synchronous serial data protocol that enables the microprocessor or microcontroller and peripheral devices to communicate with each other. The SPI controller provides a serial interface compliant with the Motorola SPI, Texas Instruments synchronous serial, and National Semiconductor MICROWIRE™ formats. In addition, SPI supports interfacing with large SPI Flash and EEPROM devices and a hardware-based slave protocol engine. There are two identical SPI controllers in the PIC64GX1000. The SPI peripherals support the following features:

- Master and Slave modes
- Selectable slaves (up to 8)
- Configurable slave select operation
- Configurable clock polarity
- Configurable as GPIO if not used
- Separate transmit (Tx) and receive (Rx) FIFOs to reduce interrupt service loading

## 7.9 Multi-Mode UART

The PIC64GX1000 multi-mode universal asynchronous/synchronous receiver/transmitter (MMUART) performs serial-to parallel conversion on data originating from modems or other serial devices, and performs parallel-to serial conversion on data from the CPUs. PIC64GX1000 contains 5 identical MMUARTs. The MMUART is software compatible with the popular 16550 UART device.

The MMUART peripherals support the following features:

- Asynchronous and synchronous operations
- Full programmable serial interface characteristics
- Data width is programmable to 5, 6, 7, or 8 bits
- Even, odd, or no-parity bit generation/detection
- 1, 1½, and 2 stop bit generation
- 9-bit address flag capability used for multi-drop addressing topologies
- Separate transmit (Tx) and receive (Rx) FIFOs to reduce processor interrupt service loading
- Single-wire half-duplex mode in which Tx pad can be used for bi-directional data transfer
- Local interconnect network (LIN) header detection and auto baud rate calculation
- Communication with ISO 7816 smart cards
- Fractional baud rate capability

- Return to Zero Inverted (RZI) mod/demod blocks that allow infrared data association (IrDA) and serial infrared (SIR) communications
- The most significant bit (MSB) or the least significant bit (LSB) as the first bit while sending or receiving data
- Configurable as GPIO if not used as an MMUART

### 7.9.1 I<sup>2</sup>C

Philips inter-integrated circuit (I<sup>2</sup>C) is a two-wire serial bus interface that provides data transfer between many devices. PIC64GX1000 contains two identical I<sup>2</sup>C peripherals that provide a mechanism for serial communication between the PIC64GX1000 and external I<sup>2</sup>C compliant devices. I<sup>2</sup>C peripherals support the following features:

- Master and Slave modes
- 7-bit addressing format and data transfers up to 100 Kbps in Standard mode and up to 400 Kbps in Fast mode
- Multi-master collision detection and arbitration
- Own slave address and general call address detection
- Second slave address detection
- System management bus (SMBus) timeout and real-time idle condition counters
- Input glitch or spike filters
  - Configurable as GPIO if not used as an MMUART

### 7.10 Real Time Counter

PIC64GX1000 integrates a real time calendar with built-in clock prescaler and an alarm wake-up comparator. It has the following two modes of operation:

- Real time Calendar (counts seconds, minutes, hours, days, weeks, months, and years)
- Binary counter (counts from 0 to  $2^{43}$ )

### 7.11 Watchdog Timer

There are five watchdogs, one per CPU. The watchdog guards against system crashes by requiring that it is regularly serviced by the assigned processor. The watchdog is not enabled at reset and generates a NMI upon reaching the trigger value. Once enabled, the watchdog cannot be disabled.

### 7.12 Timer

The PIC64GX1000 system timer consists of two programmable 32-bit decrement counters that generate interrupts to the core complex. The two 32-bit timers are identical and have the following features:

- One-shot mode
- Periodic mode
- Concatenation mode in which two 32-bit timers can be concatenated to create a 64-bit timer
- Option to enable or disable the interrupt requests when timer reaches zero
- Controls to start, stop, and reset the timer

### 7.13 Peripheral Memory SECDED Reporting and Error Injection

The Gigabit Ethernet MACs, the MMC 5.1 controller, the USB OTG controller, the CAN controllers, the crypto core and memory in the built-in Core Complex DDR controller are protected with by single-error correct, double-error detect (SECDED) error correction code (ECC) subsystem, which adds 7 bits to 32-bit memories and 8 bits to 64-bit memories. The memories within the CPU system have their own ECC control and reporting systems. The external DDR memory supports optional ECC. Each Core Complex internal memory system has its own set of control and status registers,

which consists of a status, interrupt enable, count, and error injection registers. When a two-bit error is detected, the data will not be corrected. If the data is being read over an internal AMBA bus, the system will respond with an APB, AHB, or AXI error response marking the data as corrupted. If the interrupt is enabled, an interrupt will also be generated. If the data is not being read over an AMBA bus like USB, CAN, Ethernet transmit data then corrupted data will be transmitted, and an interrupt generated. The system can then respond to the bus error event or interrupt and take the appropriate recovery action. ECC error injection is supported to ease customer validation of the error correcting subsystem. Data may be written with 1-, 2-, or 3-bit errors by setting the appropriate EDAC error injection control registers.

## 7.14 DMA Controller

The PIC64GX1000 Core Complex Direct Memory Access (DMA) controller supports up to 4 channels of independent simultaneous transfers. Each channel has its own set of control registers and two interrupts, complete and error. Bus transaction sizes are programmable and the transactions can be auto-loaded into the DMA engine. The DMA engine works in conjunction with hart software services (firmware running on the E51).

## 8. Boot Process

The Security Controller controls the start up of the CPU Core Complex harts based on the selected boot mode:

- No-Boot is used by blank devices or when debugging embedded software where code should not be executed on power up
- Standard Boot is used where the Microprocessor Subsystem harts start executing non-secured code from eNVM on power up
- Secure Boot implements Microchip-supplied factory-secure boot authentication of the eNVM content

## 9. PCI Express

Each PIC64GX1000 integrates two low-power, built-in PCIe Gen2 controllers, allowing seamless and easy connectivity to one or more host processors. The two PCIe controllers are shared across two quads.

### 9.1 PCI Express Features

The following are PCIe features:

- ×1, ×2, and ×4 lane support (FCS package only supports ×1)
- Suitable for root port
- PCI Express base specification revision 2.0 and 1.1 compliant
- Single function capability
- Advanced error reporting (AER) support
- Integrated clock domain crossing (CDC) to support user-selected AXI4 frequency
- Lane reversal support
- Legacy PCI power management support
- Native active state power management L0s and L1 state support
- Power management event (PME message)
- MSI and legacy INT message support
- Latency tolerance reporting (LTR)
- L1 PM sub-states with CLKREQ
- Address translation tables between the PCIe and AXI4 domains

#### 9.1.1 PCI Express DMA Engines

Each PCIe controller supports the following built-in DMA modes:

- Two DMA channels
- Eight outstanding read and write requests
- Completion reordering support
- Flexible scatter-gather DMA modes, including dynamic DMA control per descriptor
- Optional DMA engine reporting to the descriptor to ease software management
- Fetching of up to three descriptors to optimize throughput

## 10. Tools

The following is a list of tools compatible with PIC64GX1000:

### IDE

- MPLAB® with RV64

### Linux OS/BSP

- Linux build systems
  - YoctoBSP: Yocto-based Linux build system
- Linux4Microchip
- Commercial-Canonical – Ubuntu

### RTOS

- OpenSource – FreeRTOS, Zephyr, RTEMS, NuttX, SafeRTOS
- Commercial - Siemens Nucleus, Sel4 uKernel, PikeOS, WindRiver - VXWorks, Greenhills - µvelOSity & Integrity

### GitHub

- Bare-metal drivers, start-up code, RTOS, Linux build systems, BSP and reference examples
- For more information, see [github.com/pic64gx](https://github.com/pic64gx)

### Mi-V Ecosystem

The Mi-V RISC-V ecosystem is a continuously expanding, comprehensive suite of tools and design resources developed by Microchip and numerous third parties to fully support RISC-V designs. The Mi-V ecosystem aims to increase adoption of RISC-V ISA and Microchip's RISC-V compute products.

For more information on tools and kits, see the product page at [www.microchip.com/en-us/products/microprocessors/64-bit-mpus/pic64gx](https://www.microchip.com/en-us/products/microprocessors/64-bit-mpus/pic64gx).

## 11. Compatibility with PolarFire SoC

The PIC64GX1000 inherits the RISC-V based core complex, bus structures, security, and memory subsystems from the PolarFire SoC. In applications where additional flexibility, capability, and expansion are necessary, the PolarFire SoC, as an FPGA harnesses the processing power of the PIC64GX1000 and a large programmable fabric. The PIC64GX1000 is directly pin-compatible with the PolarFire SoC.

## 12. Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the most current publication.

Revision	Date	Description
B	10/2024	<p>The following changes were made in this revision:</p> <ul style="list-style-type: none"><li>Updated operating frequencies in <a href="#">Features</a>.</li><li>Updated block diagram in <a href="#">Overview</a>.</li><li>Updated number of lanes for the PCIe Gen2 Host for FCS variant and MIPI CSI-2 values <a href="#">Product Family Overview</a>.</li><li>Added that FCS package only supports ×1 in <a href="#">PCI Express Features</a>.</li><li>Renamed Processor I/O section to <a href="#">Peripherals</a>. Changed HDMI2.0 (FCV version) to HDMI 1.4. Added PCIe Gen2 and 1G Ethernet x2.</li><li>Removed DDR3 memory SECDED capability in <a href="#">DDR Memory Controller</a>.</li><li>Removed MII mode support from <a href="#">PHY Interfaces</a>.</li><li>Added link to GitHub PIC64-GX site in <a href="#">Tools</a>.</li></ul>
A	07/2024	Initial revision

## Microchip Information

### The Microchip Website

Microchip provides online support via our website at [www.microchip.com/](http://www.microchip.com/). This website is used to make files and information easily available to customers. Some of the content available includes:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQs), technical support requests, online discussion groups, Microchip design partner program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

### Product Change Notification Service

Microchip's product change notification service helps keep customers current on Microchip products. Subscribers will receive email notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, go to [www.microchip.com/pcn](http://www.microchip.com/pcn) and follow the registration instructions.

### Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Embedded Solutions Engineer (ESE)
- Technical Support

Customers should contact their distributor, representative or ESE for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in this document.

Technical support is available through the website at: [www.microchip.com/support](http://www.microchip.com/support)

### Microchip Devices Code Protection Feature

Note the following details of the code protection feature on Microchip products:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner, within operating specifications, and under normal conditions.
- Microchip values and aggressively protects its intellectual property rights. Attempts to breach the code protection features of Microchip product is strictly prohibited and may violate the Digital Millennium Copyright Act.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is "unbreakable". Code protection is constantly evolving. Microchip is committed to continuously improving the code protection features of our products.

### Legal Notice

This publication and the information herein may be used only with Microchip products, including to design, test, and integrate Microchip products with your application. Use of this information in any other manner violates these terms. Information regarding device applications is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure

that your application meets with your specifications. Contact your local Microchip sales office for additional support or, obtain additional support at [www.microchip.com/en-us/support/design-help/client-support-services](http://www.microchip.com/en-us/support/design-help/client-support-services).

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, OR CONSEQUENTIAL LOSS, DAMAGE, COST, OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION.

Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

## Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, CryptoMemory, CryptoRF, dsPIC, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Kleer, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, ClockWorks, The Embedded Control Solutions Company, EtherSynch, Flashtec, Hyper Speed Control, HyperLight Load, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, TimeCesium, TimeHub, TimePictra, TimeProvider, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, Clockstudio, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, EyeOpen, GridTime, IdealBridge, IGaT, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, IntelliMOS, Inter-Chip Connectivity, JitterBlocker, Knob-on-Display, MarginLink, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, mSiC, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, Power MOS IV, Power MOS 7, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SmartHLS, SMART-I.S., storClad, SQL, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, Trusted Time, TSHARC, Turing, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2024, Microchip Technology Incorporated and its subsidiaries. All Rights Reserved.

ISBN:

## **Quality Management System**

For information regarding Microchip's Quality Management Systems, please visit  
[www.microchip.com/quality](http://www.microchip.com/quality).

# Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
<b>Corporate Office</b> 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Fax: 480-792-7277 Technical Support: <a href="http://www.microchip.com/support">www.microchip.com/support</a> Web Address: <a href="http://www.microchip.com">www.microchip.com</a>	<b>Australia - Sydney</b> Tel: 61-2-9868-6733 <b>China - Beijing</b> Tel: 86-10-8569-7000 <b>China - Chengdu</b> Tel: 86-28-8665-5511 <b>China - Chongqing</b> Tel: 86-23-8980-9588 <b>China - Dongguan</b> Tel: 86-769-8702-9880 <b>China - Guangzhou</b> Tel: 86-20-8755-8029 <b>China - Hangzhou</b> Tel: 86-571-8792-8115 <b>China - Hong Kong SAR</b> Tel: 852-2943-5100 <b>China - Nanjing</b> Tel: 86-25-8473-2460 <b>China - Qingdao</b> Tel: 86-532-8502-7355 <b>China - Shanghai</b> Tel: 86-21-3326-8000 <b>China - Shenyang</b> Tel: 86-24-2334-2829 <b>China - Shenzhen</b> Tel: 86-755-8864-2200 <b>China - Suzhou</b> Tel: 86-186-6233-1526 <b>China - Wuhan</b> Tel: 86-27-5980-5300 <b>China - Xian</b> Tel: 86-29-8833-7252 <b>China - Xiamen</b> Tel: 86-592-2388138 <b>China - Zhuhai</b> Tel: 86-756-3210040	<b>India - Bangalore</b> Tel: 91-80-3090-4444 <b>India - New Delhi</b> Tel: 91-11-4160-8631 <b>India - Pune</b> Tel: 91-20-4121-0141 <b>Japan - Osaka</b> Tel: 81-6-6152-7160 <b>Japan - Tokyo</b> Tel: 81-3-6880-3770 <b>Korea - Daegu</b> Tel: 82-53-744-4301 <b>Korea - Seoul</b> Tel: 82-2-554-7200 <b>Malaysia - Kuala Lumpur</b> Tel: 60-3-7651-7906 <b>Malaysia - Penang</b> Tel: 60-4-227-8870 <b>Philippines - Manila</b> Tel: 63-2-634-9065 <b>Singapore</b> Tel: 65-6334-8870 <b>Taiwan - Hsin Chu</b> Tel: 886-3-577-8366 <b>Taiwan - Kaohsiung</b> Tel: 886-7-213-7830 <b>Taiwan - Taipei</b> Tel: 886-2-2508-8600 <b>Thailand - Bangkok</b> Tel: 66-2-694-1351 <b>Vietnam - Ho Chi Minh</b> Tel: 84-28-5448-2100	<b>Austria - Wels</b> Tel: 43-7242-2244-39 Fax: 43-7242-2244-393 <b>Denmark - Copenhagen</b> Tel: 45-4485-5910 Fax: 45-4485-2829 <b>Finland - Espoo</b> Tel: 358-9-4520-820 <b>France - Paris</b> Tel: 33-1-69-53-63-20 Fax: 33-1-69-30-90-79 <b>Germany - Garching</b> Tel: 49-8931-9700 <b>Germany - Haan</b> Tel: 49-2129-3766400 <b>Germany - Heilbronn</b> Tel: 49-7131-72400 <b>Germany - Karlsruhe</b> Tel: 49-721-625370 <b>Germany - Munich</b> Tel: 49-89-627-144-0 Fax: 49-89-627-144-44 <b>Germany - Rosenheim</b> Tel: 49-8031-354-560 <b>Israel - Hod Hasharon</b> Tel: 972-9-775-5100 <b>Italy - Milan</b> Tel: 39-0331-742611 Fax: 39-0331-466781 <b>Italy - Padova</b> Tel: 39-049-7625286 <b>Netherlands - Drunen</b> Tel: 31-416-690399 Fax: 31-416-690340 <b>Norway - Trondheim</b> Tel: 47-72884388 <b>Poland - Warsaw</b> Tel: 48-22-3325737 <b>Romania - Bucharest</b> Tel: 40-21-407-87-50 <b>Spain - Madrid</b> Tel: 34-91-708-08-90 Fax: 34-91-708-08-91 <b>Sweden - Gothenberg</b> Tel: 46-31-704-60-40 <b>Sweden - Stockholm</b> Tel: 46-8-5090-4654 <b>UK - Wokingham</b> Tel: 44-118-921-5800 Fax: 44-118-921-5820
<b>Atlanta</b> Duluth, GA Tel: 678-957-9614 Fax: 678-957-1455 <b>Austin, TX</b> Tel: 512-257-3370 <b>Boston</b> Westborough, MA Tel: 774-760-0087 Fax: 774-760-0088 <b>Chicago</b> Itasca, IL Tel: 630-285-0071 Fax: 630-285-0075 <b>Dallas</b> Addison, TX Tel: 972-818-7423 Fax: 972-818-2924 <b>Detroit</b> Novi, MI Tel: 248-848-4000 <b>Houston, TX</b> Tel: 281-894-5983 <b>Indianapolis</b> Noblesville, IN Tel: 317-773-8323 Fax: 317-773-5453 Tel: 317-536-2380 <b>Los Angeles</b> Mission Viejo, CA Tel: 949-462-9523 Fax: 949-462-9608 Tel: 951-273-7800 <b>Raleigh, NC</b> Tel: 919-844-7510 <b>New York, NY</b> Tel: 631-435-6000 <b>San Jose, CA</b> Tel: 408-735-9110 Tel: 408-436-4270 <b>Canada - Toronto</b> Tel: 905-695-1980 Fax: 905-695-2078			