AMD64 Zen6 Platform Quality of Service (PQOS) Extensions

Publication # **69193**	Revision: **1.00**
Issue Date: **March 2026**	

Future AMD APM releases that describe this feature supersede the information in this document

Revision History

Date	Revision	Change Description
March 2026	1.00	Initial release

AMD Zen 6 PQOS Extensions

This document describes new Platform Quality of Service (PQOS) features available on some Zen6 products.

- Global Bandwidth Enforcement (GLBE)
- Global Slow Bandwidth Enforcement (GLSBE)
- Privilege-Level Zero Association (PLZA)

Refer to *"AMD64 Architecture Programmer's Manual"*, Volume 2, Section 19 *"Platform Quality of Service (PQOS) Extension"* for a description of the full PQOS feature set.

Global Bandwidth Enforcement (GLBE)

Processors which set CPUID Fn8000_00020_EBX_x0[GLBE] (bit 7) = 1 support Global Bandwidth Enforcement (GLBE). GLBE provides a mechanism for software to specify bandwidth limits for groups of logical processors that span multiple QOS Domains. This collection of QOS Domains is referred to as the GLBE Control Domain. The GLBE Ceiling is a bandwidth ceiling for L3 External Bandwidth competitively shared between all logical processors in a given Class of Service (COS) across all QOS Domains within the GLBE Control Domain. This complements L3 External Bandwidth Enforcement (L3BE) which provides L3 External Bandwidth control on a per QOS Domain granularity.

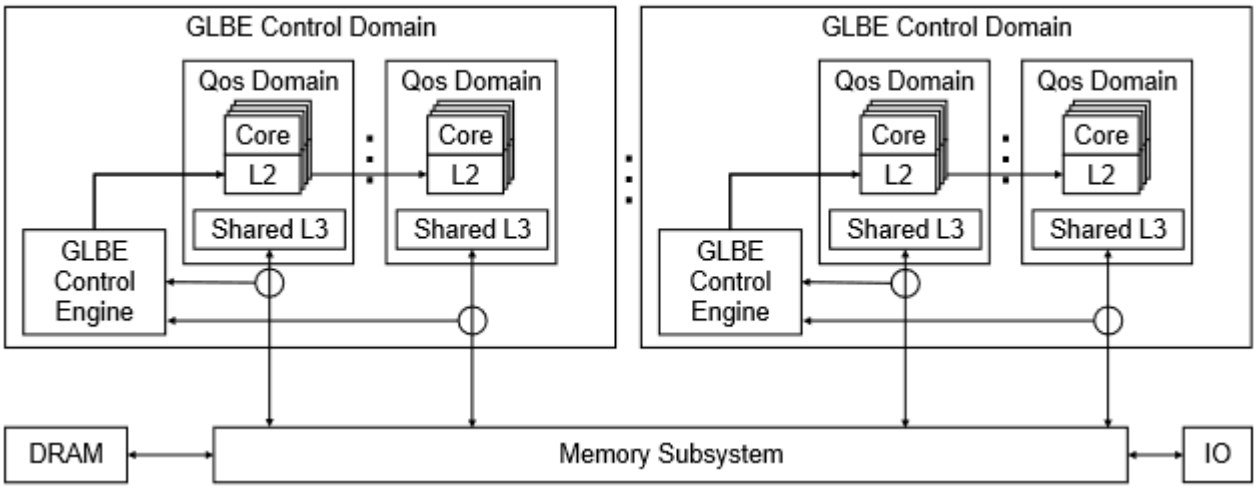

By default, all QOS Domains in the system are included in a single GLBE Control Domain. However, BIOS options may establish several GLBE Control Domains within the system. More information can be found in the *Processor Programming Reference Manual (PPR)* applicable to your product.

The attributes and capabilities of GLBE are enumerated by CPUID Fn8000_0020_x7.

CPUID Fn8000_0020_EAX_x7 Global Memory Bandwidth Enforcement (ECX=7)

Bits	Field Name	Description
31:0	BW_LEN	Global Memory Bandwidth Enforcement Ceiling Width. This field determines the size of the Ceiling field of the L3QOS_GL_BW_CONTROL_n MSRs.

CPUID Fn8000_0020_EBX_x7 Global Memory Bandwidth Enforcement (ECX=7)

Bits	Field Name	Description
15:0	BW_MULT	Identifies the Units of the bandwidth ceiling for GLBE as configured by the Ceiling field of the L3QOS_GL_BW_CONTROL_n MSRs. Units = 1/8 GB/s * (BW_MULT+1)

CPUID Fn8000_0020_EDX_x7 Global Memory Bandwidth Enforcement (ECX=7)

Bits	Field Name	Description
31:0	COS_MAX	Maximum COS number for Global Memory Bandwidth Enforcement.

GLBE MSRs

Bandwidth limits are specified per COS using the L3QOS_GL_BW_CONTROL_n MSRs where n is the COS. This set of registers starts at MSR address C000_0600 and continues through C000_0600 + (COS_MAX). The format of the L3QOS_GL_BW_CONTROL_n MSRs is shown below. There is one L3QOS_GL_BW_CONTROL_n register for each COS in each QOS Domain. The maximum COS number supported by GLBE is identified by CPUID Fn8000_0020_EDXx7[COS_MAX] (bits 31:0).

L3 Global Bandwidth Control Register (L3QOS_GL_BW_CONTROL_n) (MSR C000_0600h+n)

Bits	Mnemonic	Description	Access Type
63:BW_LEN+1	-	Reserved	MBZ
BW_LEN	U	Unlimited Bandwidth	R/W
BW_LEN-1:0	CEILING	Bandwidth Ceiling	R/W

The fields within L3QOS_GL_BW_CONTROL_n register are

- *U (Unlimited Bandwidth)* – bit BW_LEN. When cleared, the L3 External Bandwidth of logical processors belonging to this COS counts towards the bandwidth limit configured with CEILING and is limited by GLBE. When set, the L3 External Bandwidth of logical processors belonging to this COS does not count towards the bandwidth limit and is not limited by GLBE. At reset, the U bit for all L3QOS_GL_BW_CONTROL_n MSRs is set. This field may be set to different values in each QOS Domain within a single GLBE Control Domain.
- *CEILING (Bandwidth Ceiling)* – bits BW_LEN-1:0. Specifies a limit on the combined L3 External Bandwidth consumed by logical processors associated with this COS across all QOS Domains in the GLBE Control Domain. The limit is expressed in units reported by CPUID Fn8000_0020_EBX_x7[BW_MULT] (bits 15:0). This value must be configured to the same value for a given COS in all QOS Domains in the GLBE Control Domain for consistent operation.

Using GLBE

To apply a GLBE limit to a set of logical processors system software should associate all logical processors within that set with the same COS using PQR_ASSOC MSRs as described in *"AMD64 Architecture Programmer's Manual"*, Volume 2, Section 19.3.1.1 *"PQR_ASSOC"*. Then, system software uses L3QOS_GL_BW_CONTROL_n MSRs to specify the intended ceiling of combined L3 External Bandwidth for this COS and clears the U (Unlimited Bandwidth) bit. The same CEILING value must be configured for a given COS in each QOS Domain within the GLBE Control Domain. Using different CEILING values in different QOS Domains may lead to inconsistent GLBE bandwidth control. The limit specified in CEILING is an absolute value specified in the units reported by CPUID Fn8000_0020_EBX_x7[BW_MULT] (bits 15:0). The value programmed in this register may exceed the total bandwidth which the system can provide to all logical processors.

At reset, the U bit of all L3QOS_GL_BW_CONTROL MSRs is set to 1, which excludes logical processors belonging to this COS in the QOS Domain from GLBE. System software may clear the U bit and thus enable GLBE for only a subset of QOS Domains within the GLBE Control Domain.

For example, consider a system containing 4 QOS Domains in a single GLBE Control Domain. All the logical processors in all the QOS Domains are associated with COS 0. L3QOS_GL_BW_CONTROL_0[CEILING] is set to 100 GB/s on all QOS Domains. L3QOS_GL_BW_CONTROL_0[U] is set to 0 on QOS Domains 0-2, and to 1 on QOS Domain 3. COS 0 in QOS Domains 0-2 together will be limited to a total of 100 GB/s of L3 External Bandwidth. COS 0 in QOS Domain 3 will not be limited by GLBE and the L3 External Bandwidth it consumes will not be counted towards the 100 GB/s GLBE ceiling.

The GLBE Ceiling does not ensure that the specified bandwidth will be available to the logical processors associated with a given COS. Cases where the logical processors may not be able to reach their allocated bandwidth include (but are not limited to):

- Maximum system bandwidth may be less than the specified limit.
- The sum of the ceilings for all COS within a GLBE Control Domain may exceed the total bandwidth the system can supply to that GLBE Control Domain.
- I/O or other system entities may consume system bandwidth and result in less bandwidth being available to the GLBE Control Domain.
- Large amounts of write traffic may affect the memory system's ability to deliver read bandwidth.
- Other PQOS Bandwidth Enforcement Features (L3BE, L3SMBE, GLSBE) may limit bandwidth available to the GLBE Control Domain.

GLBE measures L3 External Bandwidth, not DRAM bandwidth. Reasons that bandwidth measured by GLBE may differ from System Memory Bandwidth include (but are not limited to):

- Bandwidth measured by GLBE may originate from a different QOS Domain and not from system memory.
- Bandwidth measured by GLBE may include cache coherency traffic.
- I/O or other system entities outside of a GLBE Control Domain may access system memory without being measured or controlled by a GLBE Control Domain.

CDP Interaction with GLBE

Similar to L3 Bandwidth Enforcement (L3BE), the mapping of a COS to its associated L3QOS_GL_BW_CONTROL register is changed if Code and Data Prioritization (CDP) is enabled. GLBE

does not support different bandwidth limits for data vs. code access. When CDP is enabled, for a given COS, the bandwidth limits for data and code access are specified by MSR C000_0600h + (2*COS) and bandwidth limits specified by MSR C000_0601h + (2*COS) have no effect. For example, if CDP is enabled, the bandwidth limit for data access by logical processors associated with COS 3 are specified in L3QOS_GL_BW_CONTROL_6. Note that enabling CDP reduces the effective number of unique COS values by half. Specifying a COS outside of the valid range will result in undefined behavior.

Global Slow Bandwidth Enforcement (GLSBE)

Processors which set CPUID Fn8000_00020_EBX_x0[GLSBE] (bit 8) = 1 support AMD PQOS Global Slow Bandwidth Enforcement (GLSBE). GLSBE is a feature similar to GLBE (Global Bandwidth Enforcement) which provides a mechanism for software to specify bandwidth limits for groups of logical processors that span multiple QOS Domains which access memory designated as Slow Memory by the system. The collection of QOS Domains for GLSBE is identical to that for GLBE and referred to as the GLBE Control Domain. The GLSBE ceiling is a bandwidth ceiling for L3 External Slow Memory Bandwidth which is competitively shared between all logical processors in a given COS across all QOS Domains within the GLBE Control Domain. This complements L3 Slow Memory Bandwidth Enforcement (L3SMBE) which provides L3 External Slow Memory Bandwidth control on a per QOS Domain granularity.

The attributes and capabilities of GLSBE are enumerated by CPUID Fn8000_0020_x8.

CPUID Fn8000_0020_EAX_x8 Global Slow Memory Bandwidth Enforcement (ECX=8)

Bits	Field Name	Description
31:0	BW_LEN	Global Slow Memory Bandwidth Enforcement Ceiling Width. This field determines the size of the Ceiling field of the L3QOS_GL_SLOWBW_CONTROL_n MSRs.

CPUID Fn8000_0020_EBX_x8 Global Slow Memory Bandwidth Enforcement (ECX=8)

Bits	Field Name	Description
15:0	BW_MULT	Identifies the Units of the Ceiling for GLSBE as configured with the Ceiling field of the L3QOS_GL_SLOWBW_CONTROL_n MSRs. Units = 1/8 GB/s * (BW_MULT+1).

CPUID Fn8000_0020_EDX_x8 Global Slow Memory Bandwidth Enforcement (ECX=8)

Bits	Field Name	Description
31:0	COS_MAX	Maximum COS number for Global Slow Memory Bandwidth Enforcement.

GLSBE MSRs

Bandwidth limits are specified per COS using the L3QOS_GL_SLOWBW_CONTROL_n MSRs where n is the COS. This set of registers starts at MSR address C000_0680h and continues through C000_0680h + (COS_MAX). The format of the L3QOS_GL_SLOWBW_CONTROL_n MSRs is shown below. There is one

L3QOS_GL_SLOWBW_CONTROL_n register for each COS in each QOS Domain. The maximum COS number supported by GLSBE is identified by CPUID Fn8000_0020_EDXx8[COS_MAX] (bits 31:0).

L3 Global Slow Bandwidth Control Register (L3QOS_GL_SLOWBW_CONTROL_n) (MSR C000_0680h+n)

Bits	Mnemonic	Description	Access Type
63:BW_LEN+1	-	Reserved	MBZ
BW_LEN	U	Unlimited Bandwidth	R/W
BW_LEN-1:0	CEILING	Bandwidth Ceiling	R/W

The fields within L3QOS_GL_SLOWBW_CONTROL_n register are

- *U (Unlimited Bandwidth)* – bit BW_LEN. When cleared, the L3 External Slow Memory Bandwidth of logical processors belonging to this COS counts towards the bandwidth limit configured with CEILING and is limited by GLSBE. When set, the L3 External Slow Memory Bandwidth of logical processors belonging to this COS does not count towards the bandwidth limit and is not limited by GLSBE. At reset, the U bit for all L3QOS_GL_SLOWBW_CONTROL_n MSRs is set. This field may be set to different values in each QOS Domain within a single GLBE Control Domain.
- *CEILING (Bandwidth Ceiling)* – bits BW_LEN-1:0. Specifies a limit on the combined L3 External Slow Memory Bandwidth consumed by logical processors associated with this COS across all QoS Domains in the GLBE Control Domain. The limit is expressed in units reported by CPUID Fn8000_0020_EBX_x8[BW_MULT] (bits 15:0). This value must be configured to the same value for a given COS in all QoS Domains in the GLBE Control Domain for consistent operation.

Using GLSBE

Usage of GLSBE is identical to GLBE usage except that it is enumerated by CPUID Fn8000_0020_*x8 and controlled by L3QOS_GL_SLOWBW_CONTROL_n MSRs. Refer to the “Using GLBE” section for more details.

The GLSBE Ceiling is an upper bound on the combined L3 External Bandwidth to Slow Memory of all the logical processors in a COS within the QOS Domains included in the GLBE Control Domain. The GLSBE Ceiling does not ensure that the specified bandwidth will be available to the logical processors running with a given COS. For a list of cases where logical processors in a given COS may not be able to reach their allocated bandwidth and how bandwidth measured by GLSBE can be different from system Slow Memory bandwidth, refer to the lists of cases shown in the “Using GLBE” section.

CDP Interaction with GLSBE

Similar to L3 Bandwidth Enforcement (L3BE), the mapping of a COS to its associated L3QOS_GL_SLOWBW_CONTROL register is changed if Code and Data Prioritization (CDP) is enabled. GLSBE does not support different bandwidth limits for data vs. code access. When CDP is enabled, for a given COS, the bandwidth limits for data and code access are specified by MSR C000_0680h + (2*COS) and bandwidth limits specified by MSR C000_0681h + (2*COS) have no effect. For example, if CDP is enabled, the bandwidth limit for data access by logical processors associated with COS 3 are specified in L3QOS_GL_SLOWBW_CONTROL_6. Note that enabling CDP reduces the effective number of unique COS values by half. Specifying a COS outside of the valid range will result in undefined behavior.

Privilege Level Zero Association (PLZA)

Processors which set CPUID Fn8000_0020_EBX_x0[PLZA] (bit 9) = 1 support Privilege Level Zero Association (PLZA). PLZA allows the hardware to automatically associate execution in Privilege Level Zero (CPL=0) with a specific COS and/or Resource Monitoring Identifier (RMID). The AMD PQOS feature set already has a mechanism to associate execution on each logical processor with an RMID or COS. PLZA allows the system to override this per-thread association if the logical processor is executing with CPL=0.

PLZA is controlled by the PQR_PLZA_ASSOC MSR. The format of this MSR is shown below. There is one instance of this MSR for each logical processor.

PLZA Control Register (PQR_PLZA_ASSOC) (MSR C000_03FCh)

COS_WIDTH is the width of the largest supported COS and determined as Ceil (LOG2 (1 + CPUID Fn0000_0010_EDX_x1[MAX_COS](bits 15:0))).
RMID_WIDTH is the width of the largest supported RMID and determined as Ceil (LOG2 (1 + CPUID Fn0000_000F_EBX_x0[MAX_RMID](bits 31:0))).

Bits	Mnemonic	Description	Access Type
63	PLZA_EN	PLZA Enable	R/W
62:48	-	Reserved	MBZ
47	COS_EN	PLZA Class-of-Service Association Enable	R/W
46:COS_WIDTH+32	-	Reserved	MBZ
COS_WIDTH+31:32	COS	PLZA Class-of-Service	R/W
31	RMID_EN	Resource-Monitoring-Identifier Association Enable	R/W
30:RMID_WIDTH	-	Reserved	MBZ
RMID_WIDTH-1:0	RMID	PLZA Resource-Monitoring-Identifier	R/W

The fields within PQR_PLZA_ASSOC register are

- *PLZA_EN (PLZA Enable)* – bit 63. When set, PLZA is enabled for this logical processor. Each logical processor in the QOS Domain may have a unique value for PLZA_EN. When cleared for a logical processor the COS and RMID association is specified in PQR_ASSOC MSR independent of the logical processor's privilege level.
- *COS_EN (PLZA Class-of-Service Association Enable)* – bit 47. When set and PLZA_EN is set, the Class of Service (COS) for this logical processor executing at CPL=0 is specified in the COS field instead of PQR_ASSOC[COS].
- *COS (PLZA Class-of-Service)* – bits COS_WIDTH+31:32. The COS for this logical processor executing at CPL = 0 when PLZA_EN and COS_EN are set. The size of the COS field is implementation dependent. Software must use CPUID Fn0000_0010_EDX_x1[MAX_COS](bits 15:0) to determine the largest COS value supported by the processor. An attempt to write a larger value than MAX_COS results in a #GP(0) exception.
- *RMID_EN (PLZA Resource Monitoring Identifier Association Enable)* – bit 31. When set and PLZA_EN is set, the Resource Monitoring Identifier (RMID) for this logical processor executing at CPL=0 is specified in the RMID field instead of PQR_ASSOC[RMID].

- *RMID (PLZA Resource Monitoring Identifier)* – bits RMID_WIDTH-1:0. The RMID for this logical processor executing at CPL = 0 when PLZA_EN and RMID_EN are set. The size of the RMID field is implementation dependent. Software must use CPUID Fn0000_000F_EBX_x0[MAX_RMID](bits 31:0) to determine the maximum RMID value supported by the processor. An attempt to write a larger value than MAX_RMID results in a #GP(0) exception.

PLZA Usage

Before using PLZA, the PQR_PLZA_ASSOC MSR must be configured on all logical processors in the QOS Domain. Only one COS/RMID can be associated with PLZA in a given QOS Domain and all of the fields in PQR_PLZA_ASSOC except for PLZA_EN must be set to the same value for all logical processors in the QOS Domain for consistent operation. System software must observe the following configuration sequence to ensure expected PLZA behavior.

1. Disable PLZA for all logical processors in the QOS Domain by setting PQR_PLZA_ASSOC[PLZA_EN] = 0
2. Configure the COS_EN, COS, RMID_EN, and RMID fields of PQR_PLZA_ASSOC to the desired configuration on all logical processors in the QOS Domain while PLZA remains disabled.
3. Enable PLZA for all logical processors in the QOS Domain where PLZA should be by setting PQR_PLZA_ASSOC[PLZA_EN] = 1. System software should perform this as a read-modify-write to avoid changing the value of COS_EN, COS, RMID_EN, and RMID fields of PQR_PLZA_ASSOC.

When a PLZA enabled processor is executing at CPL=0, that logical processor will use the PQOS association specified in PQR_PLZA_ASSOC instead of the RMID and COS specified in PQR_ASSOC. Transitions between PQR_PLZA_ASSOC and PQR_ASSOC, as triggered by CPL changes, are not guaranteed to be precise at CPL boundaries.

When running with SVM enabled, PQOS will use the association specified by PQR_PLZA_ASSOC if the VMM is running with CPL=0 or if the guest is running with CPL=0.

MSR Reference

MSR Address	MSR Name	Functional Group	Cross Reference
C000_03FCh	PQR_PLZA_ASSOC	Quality of Service	
C000_0600h	L3QOS_GL_BW_CONTROL_0		
C000_0600h+n	L3QOS_GL_BW_CONTROL_n		
C000_0680h	L3QOS_GL_SLOWBW_CONTROL_0		
C000_0680h+n	L3QOS_GL_SLOWBW_CONTROL_n		