



PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Process Address Space ID (PASID)
DATE:	March 31, 2011
AFFECTED DOCUMENT:	PCI Express Base Specification, Version 3.0
SPONSOR:	AMD, HP

Part I

5 **1. Summary of the Functional Changes**

This optional normative ECN defines an End-End TLP Prefix for conveying additional attributes associated with a request. The PASID TLP Prefix is an End-End TLP Prefix as defined in the *PCI Express Base Specification*. Routing elements that support End-End TLP Prefixes (i.e. have the End-End TLP Prefix Supported bit Set in the Device Capabilities 2 register) can correctly forward TLPs containing a PASID TLP Prefix.

10 The PASID TLP Prefix contains a 20 bit PASID value that, in conjunction with the Requester ID, indicates the address space of the Untranslated Address. A PASID is unique within a Requester ID (e.g. different Functions can use the same PASID value for different address spaces). The PASID TLP Prefix also contains bits to request Execute Permission and/or Privileged Mode Access permission.

15 This ECN describes the PASID TLP Prefix and usage of this prefix on Untranslated Memory Requests. A companion ECN to the *Address Translation Specification* describes usage of this prefix on Address Translation Requests and on address translation related Message Requests.

A new PASID extended capability is defined to describe the PASID Requester capabilities of an Endpoint and to Enable PASID operation. Software discovery of PASID Requester capabilities of an Endpoint is not otherwise architected.

20 No mechanism is architected to describe the PASID Completer support capabilities of an Endpoint.

No mechanism is architected to describe the PASID Requester or Completer support capabilities of a Root Complex.

25 Software may enable PASID in an Endpoint if (1) the Endpoint supports it, (2) the Completer supports it, (3) all routing elements (Switch, RC) in the path support End-End TLP Prefixes and (4) the Endpoint and the Completer both support adequate End-End TLP Prefix buffering.

PASID support is independent of support for either ATS or PRI. The Address Translation Specification describes usage of PASID TLP Prefixes in conjunction with ATS and PRI.

2. Benefits as a Result of the Changes

30 PASIDs enable sharing of a single Endpoint across multiple processes while providing each process a complete 64 bit virtual address space.

35 PASIDs enable the hierarchical management of address spaces. In Hypervisor systems, Untranslated Addresses without a PASID TLP Prefix represent Guest Physical Addresses while Untranslated Addresses with a PASID TLP Prefix represent Guest Virtual Addresses. The Hypervisor manages the Guest Physical Address to System Physical Address mappings while the Guest OS running on top of the Hypervisor manages the Guest Virtual Address to Guest Physical Address mappings.

The Execute Requested bit supports the labeling of address ranges that contain code that may be executed by an Endpoint. The Privileged Mode Requested bit supports labeling of address range protection attributes associated with two privilege levels of operation within an Endpoint (Privileged and Non-Privileged) ¹.

3. Assessment of the Impact

- 5 PASID TLP Prefix support is optional normative, and is applicable to Root Complexes, and components with Endpoints. Routing of TLPs is not affected by the presence or absence of a PASID TLP Prefix.

PASID support does not alter existing behavior. New behavior is defined for requests that contain a PASID TLP Prefix or TLPs that are the result of a request that contained a PASID TLP Prefix.

PASID support may be present with or without either ATS or PRI support.

10 4. Analysis of the Hardware Implications

Endpoints that contain the PASID Extended Capability may be enabled to generate Untranslated Memory Requests that contain the PASID TLP Prefix.

- 15 Completers that support PASID TLP Prefixes may use the PASID value to determine the affected memory location(s) and may use the Execute Requested and Privileged Mode Requested bits to perform additional access checks.

Endpoints and Completers that support the PASID TLP Prefix and also support either ATS or PRI may be enabled to generate additional TLPs with a PASID TLP Prefix. See the *Address Translation Specification* for details.

5. Analysis of the Software Implications

- 20 Config Space structure enhancements mentioned earlier enable software to discover the Function's PASID capabilities. An Endpoint is not permitted to issue requests with a PASID TLP Prefix unless software specifically enables it.

Software discovery and enablement of Completer and Root Complex PASID capabilities is outside the scope of this specification.

25

¹ These usually correspond to Supervisor and User mode operation but this correspondence is outside the scope of this specification.

Add the following item(s) to the Terms and Acronyms section:

[Process Address Space ID \(PASID\)](#)

[The Process Address Space ID, in conjunction with the Requester ID, uniquely identifies the address space associated with a transaction.](#)

Modify Section 2.1.1.1 as follows:

2.1.1.1. Memory Transactions

Memory Transactions include the following types:

- 5 Read Request/Completion
- Write Request
- AtomicOp Request/Completion

Memory Transactions use two different address formats:

- Short Address Format: 32-bit address
- 10 Long Address Format: 64-bit address

[Certain Memory Transactions can optionally have a PASID TLP Prefix containing the Process Address Space ID \(PASID\). See Section 6.20 for details.](#)

Modify Table 2-31 as follows:

Table 2-31: End-End TLP Prefix Types

End-End TLP Prefix Type	E[3:0] (b)	Description
ExtTPH	0000	Extended TPH – Refer to Section 6.17 for further details.
PASID	0001	PASID – Refer to Section 6.20 for further details.
VendPrefixE0	1110	Vendor Defined End-End TLP Prefix – Refer to Section 2.2.10.2.1 for further details.
VendPrefixE1	1111	Vendor Defined End-End TLP Prefix – Refer to Section 2.2.10.2.1 for further details.
		All other encodings are reserved.

15

Modify the ordering rules in Section 2.4.1 as follows:

2.4.1. Transaction Ordering Rules

...

A2a A Posted Request must not pass another Posted Request unless A2b applies.

20

A2b A Posted Request with RO²³ Set is permitted to pass another Posted Request.²⁴ A Posted Request with IDO Set is permitted to pass another Posted Request if the two Requester IDs are different [or if both Requests contain a PASID TLP Prefix and the two PASID values are different.](#)

...

25

B2a A Read Request must not pass a Posted Request unless B2b applies.

B2b A Read Request with IDO Set is permitted to pass a Posted Request if the two Requester IDs are different or if both Requests contain a PASID TLP Prefix and the two PASID values are different.

C2a An NPR with Data must not pass a Posted Request unless C2b applies.

5 C2b An NPR with Data and with RO Set²⁶ is permitted to pass Posted Requests. An NPR with Data and with IDO Set is permitted to pass a Posted Request if the two Requester IDs are different or if both Requests contain a PASID TLP Prefix and the two PASID values are different.

...

10 Add the following new Section 6.20:

6.20. PASID TLP Prefix

The PASID TLP Prefix is an End-End TLP Prefix as defined in Section 2.2.1. Layout of the PASID TLP Prefix is shown in Figure 6-20 and Table 6-13.

15 When a PASID TLP Prefix is present, the PASID value in the prefix, in conjunction with the requester ID, identifies the Process Address Space ID associated with the Request. Each Function has a distinct set of PASID values. PASID values used by one Function are unrelated to PASID values used by any other Function.

A PASID TLP Prefix is permitted on:

- 20 Memory Requests (including AtomicOp Requests) with Untranslated Addresses (See Section 2.2.4.1).
 Translation Requests and Translation Message Requests as defined in the *Address Translation Services Specification*.

The PASID TLP Prefix is not permitted on any other TLP.

6.20.1 Managing PASID TLP Prefix Usage

25 Usage of PASID TLP Prefixes must be specifically enabled. Unless enabled, a component is not permitted to transmit a PASID TLP Prefix.

For Endpoint Functions (including Root Complex Integrated Devices), the following rules apply:

- 30 A Function is not permitted to send and receive TLPs with a PASID TLP Prefix unless PASID Enable is Set (see Section 7.28.3):
 A Function must have a mechanism for dynamically associating use of a PASID with a particular Function context. This mechanism is device specific.
 A Function must have a mechanism to request that it gracefully stop using a specific PASID. This mechanism is device specific but must satisfy the following rules:
- A Function may support a limited number of simultaneous PASID stop requests. Software should defer issuing new stop requests until older stop requests have completed.
 - A stop request in one Function must not affect operation of any other Function.
 - A stop request must not affect operation of any other PASID within the Function.
 - A stop request must not affect operation of transactions that are not associated with a PASID.
- 35

- When the stop request mechanism indicates completion, the Function has:
 - ◆ Stopped queuing new Requests for this PASID.
 - ◆ Completed all Non-Posted Requests associated with this PASID.
 - ◆ Flushed to the host all Posted Requests addressing host memory in all TCs that were used by the PASID. The mechanism used for this is device specific (for example: a non-relaxed Posted Write to host memory or a processor read of the Function can flush TC0; a zero length read to host memory can flush non-zero TCs).
 - ◆ Optionally flushed all Peer-to-Peer Posted Requests to their destination(s). The mechanism used for this is device specific.
 - ◆ Complied with additional rules described in the *Address Translation Services Specification* if Address Translations or Page Requests were issued on the behalf of this PASID.

For Root Complexes, the following rules apply:

- A Root Complex must have a device specific mechanism for indicating support for PASID TLP Prefixes.
- A Root Complex that supports PASID TLP Prefixes must have a device specific mechanism for enabling them. By default usage of PASID TLP Prefixes is disabled.
- A Root Complex that supports PASID TLP Prefixes may optionally have a device specific mechanism for enabling them at a finer granularity than the entire Root Complex (e.g., distinct enables for a specific Root Port, Requester ID, Bus Number, Requester ID, or Requester ID/PASID combination).

6.20.2 PASID TLP Layout

A TLP may contain at most one PASID TLP Prefix.

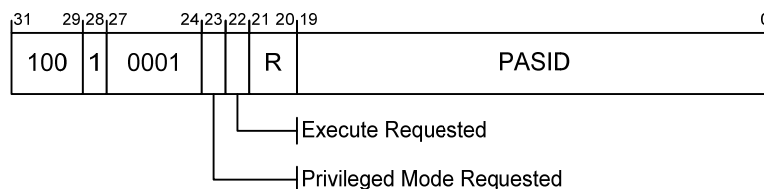


Figure 6-20: PASID TLP Prefix

Table 6-13: PASID TLP Prefix

<u>Bits</u>	<u>Columns</u>
<u>31:29</u>	<u>100b – indicating TLP Prefix</u>
<u>28</u>	<u>1b – indicating End-End TLP Prefix</u>
<u>27:24</u>	<u>0001b – indicating PASID TLP Prefix</u>

<u>Bits</u>	<u>Columns</u>
<u>23</u>	<u>Privileged Mode Requested – If Set indicates a Privileged Mode entity in the Endpoint is issuing the Request, If Clear, indicates a Non-Privileged Mode entity in the Endpoint is issuing the Request.</u> <u>Usage of this bit is specified in Section 6.20.2.3.</u>
<u>22</u>	<u>Execute Requested – If Set indicates the Endpoint is requesting Execute Permission. If Clear, indicates the Endpoint is not requesting Execute Permission.</u> <u>Usage of this bit is specified in Section 6.20.2.2.</u>
<u>21:20</u>	<u>Reserved</u>
<u>19:0</u>	<u>Process Address Space ID (PASID) – This field contains the PASID value associated with the TLP.</u> <u>Usage of this field is defined in Section 6.20.2.1.</u>

6.20.2.1. PASID field

The PASID field identifies the user process associated with a Request. This field is present in all PASID TLP Prefixes.

5 The PASID field is 20 bits wide. Endpoints and Root Complexes need not support the entire range of the field. For Endpoints, the Max PASID Width field indicates the supported range of PASID values (Section 7.28.2). For Root Complexes, an implementation specific mechanism is used to provide this information.

10 Endpoints are not permitted to send TLPs with a PASID TLP Prefix unless the PASID Enable bit (section 7.28.3) is Set. Endpoints that support the PASID TLP Prefix must signal Unsupported Request (UR) when they receive a TLP with a PASID TLP Prefix and the PASID Enable bit is Clear.

Root Complexes may optionally support TLPs with PASID TLP Prefixes. The mechanism used to detect whether a Root Complex supports the PASID TLP Prefix is implementation specific.

For Endpoints, the following rules apply:

15 The Endpoint is not permitted to send TLPs with a PASID value greater than or equal to $2^{\text{Max PASID Width}}$.

The Endpoint is optionally permitted to signal an error when it receives a TLP with a PASID value greater than or equal to $2^{\text{Max PASID Width}}$. For Requests, this error is Unsupported Request and for Completions, this error is Unexpected Completion.

For Root Complexes, the following rules apply:

20 A Root Complex is not permitted to send a TLP with a PASID value greater than it supports.

A Root Complex is optionally permitted to signal an error when it receives a TLP with a PASID value greater than it supports. For Requests, this error is Unsupported Request and for Completions, this error is Unexpected Completion.

For Completers, the following rules apply:

25 For Untranslated Memory Requests, the PASID value and the Untranslated Address are both used in determining the Translated Address used in satisfying the Request.

For address translation related TLPs, usage of this field is defined in the *Address Translation Services Specification*.

6.20.2.2. Execute Requested

If the Execute Requested bit is Set, the Endpoint is requesting permission for the Endpoint to Execute instructions in the memory range associated with this request. The meaning of Execute permission is outside the scope of this specification.

Endpoints are not permitted to send TLPs with the Execute Requested bit Set unless the Execute Permission Supported bit (section 7.28.2) and the Execute Permission Enable bit (section 7.28.3) are both Set.

For Root Complexes, the following rules apply:

- Support for Execute Requested by the Root Complex is optional. The mechanism used to determine whether a Root Complex supports Execute Requested is implementation specific.
- A Root Complex that supports the Execute Requested bit should have an implementation specific mechanism to enable it to use the bit.
- A Root Complex that supports the Execute Requested bit may have an implementation specific mechanism to enable use of the bit at a finer granularity (e.g. for a specific Root Port, for a specific Bus Number, for a specific Requester ID, or for a specific Requester ID/PASID combination), and its default value is implementation specific.

For Completers, the following rules apply:

- Completers have a concept of an effective value of the bit. For a given Request, if the Execute Requested bit is supported and its usage is enabled for the Request, the effective value of the bit is the value in the Request; otherwise the effective value of the bit is 0b.
- For Untranslated Memory Read Requests, Completers use the effective value of the bit as part of the protection check. If this protection check fails, Completers treat the Request as if the memory was not mapped.
- For Untranslated Memory Requests, other than an Untranslated Memory Read Request, the bit is Reserved.

For address translation related TLPs, usage of this bit is defined in the *Address Translation Services Specification*.

6.20.2.3. Privileged Mode Requested

If Privileged Mode Requested is Set, the Endpoint is issuing a Request that targets memory associated with Privileged Mode. If Privileged Mode Requested is Clear, the Endpoint is issuing a Request that targets memory associated with Non-Privileged Mode.

The meaning of Privileged Mode and Non-Privileged Mode and what it means for an Endpoint to be operating in Privileged or Non-Privileged Mode depends on the protection model of the system and is outside the scope of this specification.

Endpoints are not permitted to send a TLP with the Privileged Mode Requested bit Set unless both the Privileged Mode Supported bit (section 7.28.2) and the Privileged Mode Enable bit (section 7.28.3) are Set.

For Root Complexes, the following rules apply:

- ❑ Support for the Privileged Mode Requested bit by the Root Complex is optional. The mechanism used to determine whether a Root Complex supports the Privileged Mode Requested bit is implementation specific.
- ❑ A Root Complex that supports the Privileged Mode Requested bit should have an implementation specific mechanism to enable it to use the bit.
- ❑ A Root Complex that supports the Privileged Mode Requested bit may have an implementation specific mechanism to enable use of the bit at a finer granularity (e.g. for a specific Root Port, for a specific Bus Number, for a specific Requester ID, or for a specific Requester ID/PASID combination).

For Completers, the following rules apply:

- ❑ Completers have the concept of an effective value of the bit. For a given Request, if the Privileged Mode Requested bit is supported and its usage is enabled for the Request, the effective value of the bit is the value in the Request; otherwise the effective value of the bit is the 0b.
- ❑ For Untranslated Memory Requests, Completers use the effective value of the bit as part of its protection check. If this protection check fails, Completers treat the Request as if the memory was not mapped.
- ❑ For address translation related TLPs, usage of this bit is defined in the *Address Translation Services Specification*.

Add the new section 7.28:

7.28. PASID Extended Capability Structure

The presence of a PASID Extended Capability indicates that the Endpoint supports sending and receiving TLPs containing a PASID TLP Prefix. Separate support and enables are provided for the various optional features.

This capability is applicable to Endpoints and Root Complex Integrated Endpoints. For Root Ports, support and control is outside the scope of this specification.

This capability is independent of both the ATS and PRI features defined in the *Address Translation Services Specification*. Endpoints that contain a PASID Extended Capability need not support ATS or PRI. Endpoints that support ATS or PRI need not support PASID.

Figure 7-138 details allocation of the register bits in the PASID Extended Capability structure.

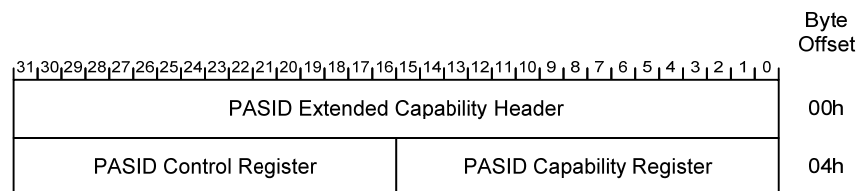


Figure 7-138: PASID Extended Capability Structure

7.28.1. PASID Extended Capability Header (Offset 00h)

Figure 7-139 details allocation of the register fields in the PASID Extended Capability header; Table 7-117 provides the respective bit definitions.

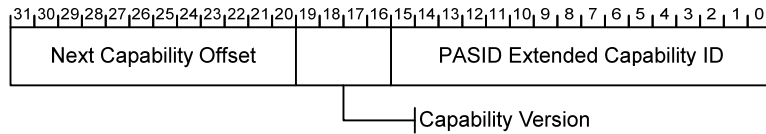


Figure 7-139: PASID Extended Capability Header

Table 7-117: PASID Extended Capability Header

<u>Bit Location</u>	<u>Register Description</u>	<u>Attributes</u>
15:0	PASID Extended Capability ID – Indicates the PASID Extended Capability structure. This field must return a Capability ID of 001Bh indicating that this is a PASID Extended Capability structure.	RO
19:16	Capability Version – This field is a PCI-SIG defined version number that indicates the version of the Capability structure present. Must be 1h for this version of the specification.	RO
31:20	Next Capability Offset – The offset to the next PCI Extended Capability structure or 000h if no other items exist in the linked list of capabilities.	RO

7.28.2. PASID Capability Register (Offset 04h)

Figure 7-140 details the allocation of register bits of the PASID Capability register; Table 7-118 provides the respective bit definitions.

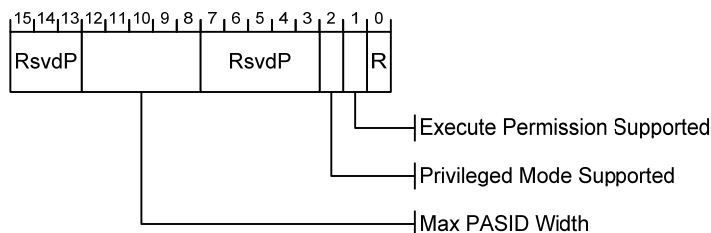


Figure 7-140: PASID Capability Register

Table 7-118: PASID Capability Register

<u>Bit Location</u>	<u>Register Description</u>	<u>Attributes</u>
<u>0</u>	<u>Reserved</u>	<u>RsvdP</u>
<u>1</u>	<u>Execute Permission Supported</u> – If Set, the Endpoint supports sending TLPs that have the Execute Requested bit Set. If Clear, the Endpoint will never Set the Execute Requested bit.	<u>RO</u>
<u>2</u>	<u>Privileged Mode Supported</u> – If Set, the Endpoint supports operating in Privileged and Non-Privileged modes, and supports sending requests that have the Privileged Mode Requested bit Set. If Clear, the Endpoint will never Set the Privileged Mode Requested bit.	<u>RO</u>
<u>7:3</u>	<u>Reserved</u>	<u>RsvdP</u>
<u>12:8</u>	<u>Max PASID Width</u> – Indicates the width of the PASID field supported by the Endpoint. The value n indicates support for PASID values 0 through 2^n-1 (inclusive). The value 0 indicates support for a single PASID (0). The value 20 indicates support for all PASID values (20 bits). This field must be between 0 and 20 (inclusive).	<u>RO</u>

7.28.3. PASID Control Register (Offset 06h)

Figure 7-141 details the allocation of register bits of the PASID Control register; Table 7-119 provides the respective bit definitions.

5

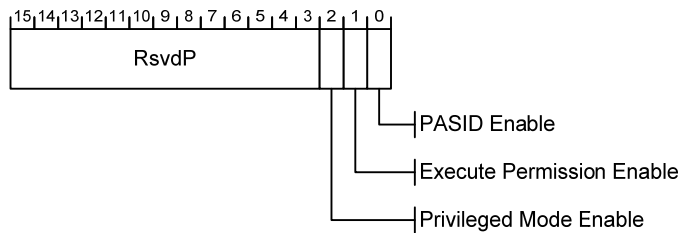


Figure 7-141: PASID Control Register

Table 7-119: PASID Control Register

<u>Bit Location</u>	<u>Register Description</u>	<u>Attributes</u>
<u>0</u>	<p>PASID Enable – If Set, the Endpoint is permitted to send and receive TLPs that contain a PASID TLP Prefix. If Clear, the Endpoint is not permitted to do so.</p> <p>Behavior is undefined if the Endpoint supports ATS and this bit changes value when ATS Enable (E) is Set (see the <i>Address Translation Specification</i>).</p> <p>Default is 0b.</p>	RW
<u>1</u>	<p>Execute Permission Enable – If Set, the Endpoint is permitted to send Requests that have Execute Requested Set. If Clear, the Endpoint is not permitted to do so.</p> <p>Behavior is undefined if the Endpoint supports ATS and this bit changes value when ATS Enable is Set (see the <i>Address Translation Specification</i>).</p> <p>If Execute Permission Supported is Clear, this bit is RsvdP.</p> <p>Default is 0b.</p>	RW / RsvdP (see description)
<u>2</u>	<p>Privileged Mode Enable – If Set, the Endpoint is permitted to send Requests that have the Privileged Mode Requested Set. If Clear, the Endpoint is not permitted to do so.</p> <p>Behavior is undefined if the Endpoint supports ATS and this bit changes value when ATS Enable is Set (see the <i>Address Translation Specification</i>).</p> <p>If Privileged Mode Supported is Clear, this bit is RsvdP.</p> <p>Default is 0b.</p>	RW / RsvdP (see description)

Change Section 3, Table 3-1 of the *PCI Code and ID Assignment Specification* as follows:

Table 3-1: Extended Capability IDs

ID	Extended Capability
001Ah	Protocol Multiplexing
<u>001Bh</u>	<u>PASID</u>

5