# PCI-SIG ENGINEERING CHANGE NOTICE

| | |
|---|---|
| **TITLE:** | PCI Express Access Control Services (ACS) |
| **DATE:** | September 27, 2005; updated October 11, 2006 |
| **AFFECTED DOCUMENT:** | PCI Express Base Specification version 1.1 |
| **SPONSOR:** | Advanced Micro Devices<br>Hewlett-Packard |

## Part I

*Changes made since the 30-day Member Review are highlighted in yellow.*

### 1. Summary of the Functional Changes

PCI Express supports multiple fabric (point-to-point and multi-level Switch) and Endpoint (single and multi-function) topologies. These topologies combined with standard PCI functionality such as hot-plug and peer-to-peer communications, present a series of access control problems for customers. Access control is required to limit what components are allowed to communicate with one another. This document proposes adding a set of access control services (ACS) to PCI Express currently not covered within the existing specifications. The proposed services are applicable to PCI Express RCs, Switches, and multi-function devices. Some of the proposed services apply to new functionality defined by the Address Translation Services (ATS) Specification, which is currently under development.

### 2. Benefits as a Result of the Changes

Customers will be able to selectively control access between PCI Express Endpoints and between Functions within a multi-function device through well-defined interfaces. The following examples illustrate some of the benefits customers will be able to enable through the use of these controls:

- ACS can be used to prevent various forms of silent data corruption by preventing PCI Express Requests from being incorrectly routed to a peer Endpoint. For example, if the address within a Request header is silently corrupted within a PCI Express Switch (such as in a store-and-forward implementation), the transaction may be incorrectly routed to a downstream Endpoint or Root Port and acted upon it as though it were a valid transaction for that component. This could result in any number of problems which may or may not be able to be detected by the associated application or service.

- ACS can be used to prevent data leakage by precluding PCI Express Requests from being routed between Functions within a multi-function device.

- ACS can be used to validate that every Request transaction between two downstream components is allowed. Validation can occur within intermediate components or within the RC itself.

- On systems where ATS is being used, ACS can be used to enable direct routing of peer-to-peer Memory Requests whose addresses have been Translated, while blocking or redirecting peer-to-peer Memory Requests whose addresses have not been Translated.

### 3. Assessment of the Impact

ACS is an optional normative capability which is applicable to RCs, Switches, and multi-function devices. ACS is implemented as a set of capabilities and control registers in the associated hardware

component. Software detects and configures ACS control registers to achieve the desired access control. Given the optional normative nature, the technology is backward compatible with existing hardware and software, i.e. if the ACS capabilities are not enabled, the hardware and software operate in compliance with the existing PCI Express 1.1 Base Specification.

The services for Switches to provide enhanced access control for ATS functionality are fully defined in this ACS ECN. Switch implementations need not wait for the completion of the ATS Specification.

### 4. Analysis of the Hardware Implications

ACS requires new ACS-capable hardware and is therefore optional normative. Hardware that is not ACS-capable will treat ACS-capable hardware per the existing PCI Express Base Specification and be fully interoperable. To utilize ACS P2P Request/Completion Redirect functionality, certain sets of components must have specific ACS capabilities. For example, to utilize ACS P2P Request Redirect with a multi-function device, the RC and any intermediate Switches must support ACS Upstream Forwarding.

When ACS is enabled, a component is required to examine each TLP to determine whether the TLP should be routed normally, blocked, or redirected. ACS is implemented as a set of capabilities and control registers. ACS-capable hardware must implement the associated capabilities and control registers to meet the desired level of access control. If ACS is not configured, then the hardware must behave per the existing PCI Express Base Specification.

For interoperability purposes, the following are supported:

- The ability to mix ACS and non-ACS capable hardware with varying degrees of access control provided. This can range from no access control to access control only at defined or selective points within a hierarchy or topology.

- A PCI Express component is allowed to implement only a subset of the access control functionality. Therefore, ACS functionality will vary by component as well as within a given component depending upon the target usage models. The extent of functionality implemented is communicated via the associated capabilities.

### 5. Analysis of the Software Implications

ACS requires new software to enable and configure ACS-capable hardware and is thus optional normative. Software that is not ACS-capable will treat ACS-capable hardware per the existing PCI Express Base Specification and be fully interoperable.

ACS functionality is accessed via a new set of ACS capabilities and control registers. ACS software accesses these structures through existing PCI Express configuration methodologies.

## Part II

### Detailed Description of the change

*Add to Terms and Acronyms*

| | |
|---|---|
| ACS | Access Control Services: A set of capabilities and control registers used to implement access control over routing within a PCI Express component. |
| ACS Violation | An error that applies to a Posted or Non-Posted Request when the Completer detects an access control violation checked by ACS. |
| P2P | Peer-to-peer. |

*Add to Chapter 2*

### 2.2.4.1. Address Based Routing Rules

- Address routing is used with Memory and I/O Requests.
- Two address formats are specified, a 64-bit format used with a 4 DW header (see Figure 2-13) and a 32-bit format used with a 3 DW header (see Figure 2-14).

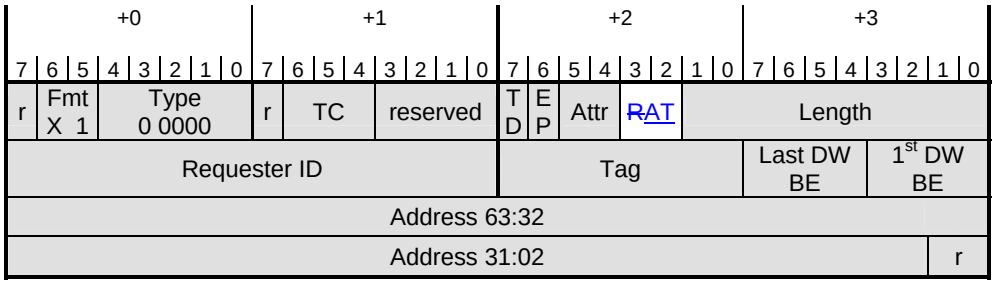


**Figure 2-13: 64-bit Address Routing**

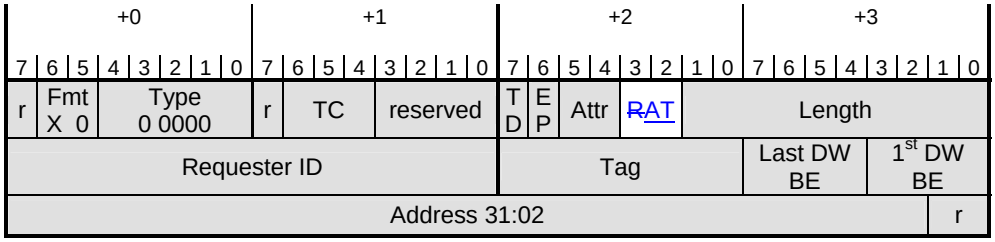


**Figure 2-14: 32-bit Address Routing**

*Note: must also update all other Base spec figures with first TLP DWORD to show the new AT field.*

- For Memory Read Requests and Memory Write Requests, the Address Type (AT) field is encoded as shown in Table 2-xx, with full descriptions contained in the Address Translation Services Specification Revision 1.0. For all other Requests, the AT field is reserved.

**Table 2-xx: Address Type (AT) Field Encodings**

| AT Coding | Basic Description |
|---|---|
| 00b | default / Untranslated |
| 01b | Translation Request |
| 10b | Translated |
| 11b | reserved |

- Address mapping to the TLP header is shown in Table 2-5.

…

*Add to Chapter 6*

#### 6.2.3.2.4.1. Completer Sending a Completion with UR/CA Status

A Completer generally sends a Completion with an Unsupported Request or Completer Abort (UR/CA) Status to signal a uncorrectable error for a Non-Posted Request.[1] If the severity of the UR/CA error[2] is non-fatal, the Completer must handle this case as an Advisory Non-Fatal Error.[3] A Completer with AER signals the non-fatal error (if enabled) by sending an ERR_COR Message. A Completer without AER sends no Error Message for this case.

…

## 6.2.7. Error Listing and Rules

…

**Table 6-4 Transaction Layer Error List**

| Error Name | Severity | Detecting Agent Action |
|---|---|---|
| … | Uncorrectable (Non-Fatal) | … |
| Unexpected Completion | | Receiver:<br>Send ERR_NONFATAL to Root Complex.<br>Log the header of the Completion that encountered the error.<br>Note that if the Unexpected Completion is a result of misrouting, the Completion Timeout mechanism will be triggered at the corresponding Requester. |
| ACS Violation | | Receiver (if checking):<br>Send ERR_NONFATAL to Root Complex.<br>Log the header of the Request that encountered the error. |
| Receiver Overflow | Uncorrectable (Fatal) | Receiver (if checking):<br>Send ERR_FATAL to Root Complex. |
| … | | … |

[1] If the Completer is returning data in a Completion, and the data is bad or suspect, the Completer is permitted to signal the error using the Error Forwarding (Data Poisoning) mechanism instead of handling it as a UR or CA.

[2] An ACS Violation error with a Non-Posted Request also results in the Completer sending a Completion with CA Status. If the severity of the ACS Violation error is non-fatal, the Completer must also handle this case as an Advisory Non-Fatal Error.

[3] If the severity is fatal, the error is not an Advisory Non-Fatal Error, and must be signaled (if enabled) with ERR_FATAL.

# 6.11. Access Control Services (ACS)

ACS defines a set of control points within a PCI Express topology to determine whether a TLP should be routed normally, blocked, or redirected. ACS is applicable to RCs, Switches, and multi-function devices[4].

ACS provides the following types of access control:

1. ACS Source Validation (V)
2. ACS Translation Blocking (B)
3. ACS P2P Request Redirect (R)
4. ACS P2P Completion Redirect (C)
5. ACS Upstream Forwarding (U)
6. ACS P2P Egress Control (E)
7. ACS Direct Translated P2P (T)

The specific requirements for each of these are discussed in the following section. The letter in parenthesis following each type is the abbreviation for the associated capability and control bits defined in Section 7.16.

ACS hardware functionality is disabled by default, and is enabled only by ACS-aware software.

## 6.11.1. ACS Component Capability Requirements

ACS functionality is reported and managed via ACS Extended Capability structures. PCI Express components are permitted to implement ACS Extended Capability structures in some, none, or all of their applicable Functions. The extent of what is implemented is communicated through capability bits in each ACS Extended Capability Structure. A given Function with an ACS Extended Capability structure may be required or forbidden to implement certain capabilities, depending upon the specific type of the Function and whether it's part of a multi-function device.

ACS is never applicable to a PCI Express to PCI Bridge Function or a Root Complex Event Collector Function, and such Functions must never implement an ACS Extended Capability Structure.

### 6.11.1.1. ACS Downstream Ports

This section applies to Root Ports and Downstream Switch Ports that implement an ACS Extended Capability structure. This section applies to Downstream Port Functions both for single-function devices and multi-function devices.

- ACS Source Validation: must be implemented.

[4] Applicable Functions within multi-function devices specifically include PCI Express Endpoints, Switch Upstream Ports, Legacy PCI Express Endpoints, and Root Complex Integrated Endpoints.

When enabled, the Downstream Port tests the Bus Number from the Requester ID of each upstream Request received by the Port to determine if it is within the Bus Number "aperture" of the Port – the inclusive range specified by the Secondary Bus Number register and the Subordinate Bus Number register.

If the Bus Number from the Requester ID of the Request is not within this aperture, this is a reported error (ACS Violation) associated with the Receiving Port (see Section 6.11.4.)

Completions are never affected by ACS Source Validation.

- ACS Translation Blocking: must be implemented.

  When enabled, the Downstream Port checks the Address Translation (AT) field of each upstream Memory Request received by the Port. If the AT field is not the default value, this is a reported error (ACS Violation) associated with the Receiving Port (see Section 6.11.4.)

  Completions are never affected by ACS Translation Blocking.

- ACS P2P Request Redirect: must be implemented by Root Ports that support peer-to-peer traffic with other Root Ports[5]; must be implemented by Switch Downstream Ports.

  ACS P2P Request Redirect is subject to interaction with the ACS P2P Egress Control and ACS Direct Translated P2P mechanisms (if implemented). See Section 6.11.3.

  When ACS P2P Request Redirect is enabled in a Switch Downstream Port, peer-to-peer Requests must be redirected upstream towards the RC.

  When ACS P2P Request Redirect is enabled in a Root Port, peer-to-peer Requests must be sent to Redirected Request Validation logic within the RC that determines whether the Request is "reflected" back downstream towards its original target, or blocked as an ACS Violation error. The algorithms and specific controls for making this determination are not architected by this specification.

  Downstream Ports never redirect Requests that are traveling downstream.

  Completions are never affected by ACS P2P Request Redirect.

- ACS P2P Completion Redirect: must be implemented by Root Ports that implement ACS P2P Request Redirect; must be implemented by Switch Downstream Ports.

  The intent of ACS P2P Completion Redirect is to avoid ordering rule violations between Completions and Requests when Requests are redirected. See Section 6.11.5.

  ACS P2P Completion Redirect does not interact with ACS controls that govern Requests.

  When ACS P2P Completion Redirect is enabled in a Switch Downstream Port, peer-to-peer Read Completions that do not have the Relaxed Ordering Attribute bit set (1b) must be redirected upstream towards the RC. Otherwise, peer-to-peer Completions must be routed normally.

---

[5] Root Port indication of ACS P2P Request Redirect or ACS P2P Completion Redirect support does not imply any particular level of peer-to-peer support by the RC, or that peer-to-peer traffic is supported at all.

When ACS P2P Completion Redirect is enabled in a Root Port, peer-to-peer Read Completions that do not have the Relaxed Ordering bit set must be handled such that they do not pass Requests that are sent to Redirected Request Validation logic within the RC. Such Completions must eventually be sent downstream towards their original peer-to-peer targets, without incurring additional ACS access control checks.

Downstream Ports never redirect Completions that are traveling downstream.

Requests are never affected by ACS P2P Completion Redirect.

- ACS Upstream Forwarding: must be implemented by Root Ports if the RC supports Redirected Request Validation; must be implemented by Switch Downstream Ports.

  When ACS Upstream Forwarding is enabled in a Switch Downstream Port, and its Ingress Port receives an upstream Request or Completion TLP targeting the Port's own Egress Port, the Port must instead forward the TLP upstream towards the RC.

  When ACS Upstream Forwarding is enabled in a Root Port, and its Ingress Port receives an upstream Request or Completion TLP that targets the Port's own Egress Port, the Port must handle the TLP as follows. For a Request, the Root Port must handle it the same as a Request that the Port "redirects" with the ACS P2P Request Redirect mechanism. For a Completion, the Root Port must handle it the same as a Completion that the Port "redirects" with the ACS P2P Completion Redirect mechanism.

  When ACS Upstream Forwarding is not enabled on a Downstream Port, and its Ingress Port receives an upstream Request or Completion TLP that targets the Port's own Egress Port, the handling of the TLP is undefined.

- ACS P2P Egress Control: implementation is optional.

  ACS P2P Egress Control is subject to interaction with the ACS P2P Request Redirect and ACS Direct Translated P2P mechanisms (if implemented). See Section 6.11.3.

  A Switch that supports ACS P2P Egress Control can be selectively configured to block peer-to-peer Requests between its Downstream Ports. Software can configure the Switch to allow none or only a subset of its Downstream Ports to send peer-to-peer Requests to other Downstream Ports. This is configured on a per Downstream Port basis.

  An RC that supports ACS P2P Egress Control can be selectively configured to block peer-to-peer Requests between its Root Ports. Software can configure the RC to allow none or only a subset of the Root Port hierarchies to send peer-to-peer Requests to other Root Port hierarchies. This is configured on a per Root Port basis.

  With ACS P2P Egress Control in Downstream Ports, controls in the Ingress Port ("sending" Port) determine if the peer-to-peer Request is blocked, and if so, the Ingress Port handles the ACS Violation error per section 6.11.4.

  Completions are never affected by ACS P2P Egress Control.

- ACS Direct Translated P2P: must be implemented by Root Ports that support Address Translation Services (ATS) and also support peer-to-peer traffic with other Root Ports[6]; must be implemented by Switch Downstream Ports.

  When ACS Direct Translated P2P is enabled in a Downstream Port, peer-to-peer Memory Requests whose Address Type (AT) field indicates a Translated address must be routed normally ("directly") to the peer Egress Port, regardless of ACS P2P Request Redirect and ACS P2P Egress Control settings. All other peer-to-peer Requests must still be subject to ACS P2P Request Redirect and ACS P2P Egress Control settings.

  Completions are never affected by ACS Direct Translated P2P.

#### 6.11.1.2. ACS Functions in Multi-Function Devices

This section applies to multi-function device ACS Functions, with the exception of Downstream Port Functions, which are covered in the preceding section.

- ACS Source Validation: must not be implemented.

- ACS Translation Blocking: must not be implemented.

- ACS P2P Request Redirect: must be implemented by Functions that support peer-to-peer traffic with other Functions.

  ACS P2P Request Redirect is subject to interaction with the ACS P2P Egress Control and ACS Direct Translated P2P mechanisms (if implemented). See Section 6.11.3.

  When ACS P2P Request Redirect is enabled in a multi-function device, peer-to-peer Requests (between Functions of the device) must be redirected upstream towards the RC.

  Completions are never affected by ACS P2P Request Redirect.

- ACS P2P Completion Redirect: must be implemented by Functions that implement ACS P2P Request Redirect.

  The intent of ACS P2P Completion Redirect is to avoid ordering rule violations between Completions and Requests when Requests are redirected. See Section 6.11.5.

  ACS P2P Completion Redirect does not interact with ACS controls that govern Requests.

  When ACS P2P Completion Redirect is enabled in a multi-function device, peer-to-peer Read Completions that do not have the Relaxed Ordering bit set must be redirected upstream towards the RC. Otherwise, peer-to-peer Completions must be routed normally.

  Requests are never affected by ACS P2P Completion Redirect.

- ACS Upstream Forwarding: must not be implemented.

---

[6] Root Port indication of ACS Direct Translated P2P support does not imply any particular level of peer-to-peer support by the RC, or that peer-to-peer traffic is supported at all.

- ACS P2P Egress Control: implementation is optional; is based on Function Numbers; controls peer-to-peer Requests between the different Functions within the multi-function device.

  ACS P2P Egress Control is subject to interaction with the ACS P2P Request Redirect and ACS Direct Translated P2P mechanisms (if implemented). See Section 6.11.3.

  Each Function within a multi-function device that supports ACS P2P Egress Control can be selectively enabled to block peer-to-peer communication with other Functions within the device. This is configured on a per Function basis. Conceptually, the Functions are interconnected through a transparent embedded switch, and access control uses logic similar to that in a Switch but validates whether the Function field within the Requester ID of a Request is allowed or not.

  With ACS P2P Egress Control in multi-function devices, controls in the "sending" Function determine if the Request is blocked, and if so, the "sending" Function handles the ACS Violation error per section 6.11.4.

  Completions are never affected by ACS P2P Egress Control.

- ACS Direct Translated P2P: must be implemented if the multi-function device Function supports Address Translation Services (ATS) and also peer-to-peer traffic with other Functions.

  When ACS Direct Translated P2P is enabled in a multi-function device Function, peer-to-peer Memory Requests whose Address Type (AT) field indicates a Translated address must be routed normally ("directly") to the peer Function, regardless of ACS P2P Request Redirect and ACS P2P Egress Control settings. All other peer-to-peer Requests must still be subject to ACS P2P Request Redirect and ACS P2P Egress Control settings.

  Completions are never affected by ACS Direct Translated P2P.

#### 6.11.1.2. Functions in Single-Function Devices

This section applies to single-function device Functions, with the exception of Downstream Port Functions, which are covered in a preceding section. No ACS capabilities are applicable, and the Function must not implement an ACS Extended Capability structure.

### 6.11.2. Interoperability

The following rules govern interoperability between ACS and non-ACS components:

- When ACS P2P Request Redirect and ACS P2P Completion Redirect are not being used, ACS and non-ACS components may be intermixed within a topology and will interoperate fully. ACS can be enabled in a subset of the ACS components without impacting interoperability.

- When ACS P2P Request Redirect, ACS P2P Completion Redirect, or both are being used, certain components in the PCI Express hierarchy must support ACS Upstream Forwarding (of upstream redirected Requests). Specifically:

- The associated Root Port[7] must support ACS Upstream Forwarding. Otherwise, how the Root Port handles upstream redirected Request or Completion TLPs is undefined. The RC must also implement Redirected Request Validation.

- Between each ACS component where P2P TLP redirection is enabled and its associated Root Port, any intermediate Switches must support ACS Upstream Forwarding. Otherwise, how such Switches handle upstream redirected TLPs is undefined.

## 6.11.3. ACS Peer-to-Peer Control Interactions

With each peer-to-peer **Request**, multiple ACS control mechanisms may interact to determine whether the Request is routed directly towards its peer-to-peer target, blocked immediately as an ACS Violation, or redirected upstream towards the RC for access validation. Peer-to-peer **Completion** redirection is determined exclusively by the ACS P2P Completion Redirect mechanism.

If ACS Direct Translated P2P is enabled in a Port/Function, peer-to-peer Memory Requests whose Address Translation (AT) field indicates a Translated address must be routed normally ("directly") to the peer Port/Function, regardless of ACS P2P Request Redirect and ACS P2P Egress Control settings. Otherwise such Requests, and unconditionally all other peer-to-peer Requests, must be subject to ACS P2P Request Redirect and ACS P2P Egress Control settings. Specifically, the applicable Egress Control Vector bit, along with the ACS P2P Egress Control Enable bit (E) and the ACS P2P Request Redirect Enable bit (R), determine how the Request is handled. See Section 7.16 for descriptions of these control bits. Table 7-xx specifies the interactions:

**Table 7-xx ACS P2P Request Redirect and ACS P2P Egress Control Interactions**

| Control Bit E | Control Bit R | Egress Control Vector Bit for the associated Egress Switch Port, Root Port, or Function | Required handling for peer-to-peer Requests |
|---|---|---|---|
| 0 | 0 | X – Don't care | Route directly to peer-to-peer target |
| 0 | 1 | X – Don't Care | Redirect upstream |
| 1 | 0 | 1 | Handle as an ACS Violation |
| 1 | 0 | 0 | Route directly to peer-to-peer target |
| 1 | 1 | 1 | Redirect upstream |
| 1 | 1 | 0 | Route directly to peer-to-peer target |

## 6.11.4. ACS Violation Error Handling

ACS Violations may occur due to either hardware or software defects / failures. To assist in fault isolation and root cause analysis, it is recommended that AER be implemented in ACS components. The AER Header Log register can log the header of the offending Request. The ACS Violation

[7] Not applicable for ACS Redirect between Functions of a multi-Function Root Complex Integrated Endpoint.

Status, Mask, and Severity bits provide positive identification of the error and increased control over error logging and signaling.

When an ACS Violation is detected, the ACS component that operates as the Completer[8] must do the following:

- For Non-Posted Requests, the Completer must generate a Completion with a Completer Abort (CA) Completion Status.
- The Completer must log and signal the ACS Violation as indicated in Figure 6-2. Note the following:
  - Even though the Completer uses a CA Completion Status when it sends a Completion, the Completer must log an ACS Violation error instead of a Completer Abort error.
  - If the severity of the ACS Violation is non-fatal and the Completer sends a Completion with CA Completion Status, this case must be handled as an Advisory Non-Fatal Error as described in Section 6.2.3.2.4.1.
- The Completer[9] must set the Signaled Target Abort bit in either its Status register or Secondary Status register as appropriate.

## 6.11.5. ACS Redirection Impacts on Ordering Rules

When ACS P2P Request Redirect is enabled, some or all peer-to-peer Requests are redirected, which can cause ordering rule violations in some cases. This section explores those cases, plus a similar case that occurs with RCs that implement "Request Retargeting" as an alternative mechanism for enforcing peer-to-peer access control.

### 6.11.5.1. Completions Passing Posted Requests

When a peer-to-peer Posted Request is redirected, a subsequent peer-to-peer non-RO[10] Read Completion that is routed directly can effectively pass the redirected Posted Request, violating the ordering rule that non-RO Read Completions must not pass Posted Requests. See Section 2.4.1.

ACS P2P Completion Redirect can be used to avoid violating this ordering rule. When ACS P2P Completion Redirect is enabled, all peer-to-peer non-RO Read Completions will be redirected, thus

---

[8] In all cases but one, the ACS component that detects the ACS Violation also operates as the Completer. The exception case is when RC Redirected Request Validation logic disallows a redirected Request. If the redirected Request came through a Root Port, that Root Port must operate as the Completer. If the redirected Request came from a Root Complex Integrated Endpoint, the associated Root Complex Event Collector must operate as the Completer.

[9] Similarly, if the Request was Non-Posted, when the Requester receives the resulting Completion with CA Completion Status, the Requester must set the Received Target Abort bit in either its Status register or Secondary Status register as appropriate.

[10] In this section, "non-RO" is an abbreviation characterizing TLPs whose Relaxed Ordering Attribute field is not set.

taking the same path as redirected peer-to-peer Posted Requests. Enabling ACS P2P Completion Redirect when some or all peer-to-peer Requests are routed directly will not cause any ordering rule violations, since it's permitted for a given Completion to be passed by any TLP other than another Completion with the same Transaction ID.

As an alternative mechanism to ACS P2P Request Redirect for enforcing peer-to-peer access control, some RCs implement "Request Retargeting", where the RC supports special address ranges for "peer-to-peer" traffic, and the RC will retarget validated upstream Requests to peer devices. Upon receiving an upstream Request targeting a special address range, the RC validates the Request, translates the address to target the appropriate peer device, and sends the Request back downstream. With retargeted Requests that are Non-posted, if the RC does not modify the Requester ID, the resulting Completions will travel "directly" peer-to-peer back to the original Requester, creating the possibility of non-RO Read Completions effectively passing retargeted Posted Requests, violating the same ordering rule as when ACS P2P Request Redirect is being used. ACS P2P Completion Redirect can be used to avoid violating this ordering rule here as well.

If ACS P2P Request Redirect and RC P2P Request Retargeting are not being used, there's no envisioned benefit to enabling ACS P2P Completion Redirect, and it's recommended not to do so because of potential performance impacts.

---

## IMPLEMENTATION NOTE

### Performance Impacts with ACS P2P Completion Redirect

While the use of ACS P2P Completion Redirect can avoid ordering violations with Completions passing Posted Requests, it also may impact performance. Specifically, all redirected Completions will have to travel up to the RC from the point of redirection and back, introducing extra latency and possibly increasing Link and RC congestion.

Since peer-to-peer Read Completions with the Relaxed Ordering bit set are never redirected (thus avoiding performance impacts), it is strongly recommended that Requesters be implemented to maximize the proper use of Relaxed Ordering, and that software enable Requesters to utilize Relaxed Ordering by setting the Enable Relaxed Ordering bit in the Device Control register.

If software enables ACS P2P Request Redirect, RC P2P Request Retargeting, or both, and software is certain that proper operation is not compromised by peer-to-peer non-RO Read Completions passing peer-to-peer[11] Posted Requests, it is recommended that software leave ACS P2P Completion Redirect disabled as a way to avoid its performance impacts.

---

[11] These include true peer-to-peer Requests that are redirected by the ACS P2P Request Redirect mechanism, as well "logically peer-to-peer" Requests routed to the RC that the RC then retargets to the peer device.

### 6.11.5.2. Requests Passing Posted Requests

When some peer-to-peer Requests are redirected but other peer-to-peer Requests are routed directly, the possibility exists of violating the ordering rules where Non-posted Requests or non-RO Posted Requests must not pass Posted Requests. See Section 2.4.1.

These ordering rule violation possibilities exist only when ACS P2P Request Redirect and ACS Direct Translated P2P are both enabled. Software should not enable both these mechanisms unless it is certain either that such ordering rule violations can't occur, or that proper operation will not be compromised if such ordering rule violations do occur.

---

## IMPLEMENTATION NOTE

### Ensuring Proper Operation with ACS Direct Translated P2P

The intent of ACS Direct Translated P2P is to optimize performance in environments where Address Translation Services (ATS) are being used with peer-to-peer communication whose access control is enforced by the RC. Permitting peer-to-peer Requests with Translated addresses to be routed directly avoids possible performance impacts associated with redirection, which introduces extra latency and may increase Link and RC congestion.

For the use model where peer-to-peer Requests with Translated addresses are permitted, but those with Untranslated addresses are to be blocked as ACS Violations, it is recommended that software enable ACS Direct Translated P2P and ACS P2P Request Redirect, and configure the Redirected Request Validation logic in the RC to block the redirected Requests with Untranslated addresses. This configuration has no ordering rule violations associated with Requests passing Posted Requests.

For the use model where some Requesters use Translated addresses exclusively with peer-to-peer Requests and some Requesters use Untranslated addresses exclusively with peer-to-peer Requests, and the two classes of Requesters don't communicate peer-to-peer with each other, proper operation is unlikely to be compromised by redirected peer-to-peer Requests (with Untranslated addresses) being passed by direct peer-to-peer Requests (with Translated addresses). It is recommended that software not enable ACS Direct Translated P2P unless software is certain that proper operation is not compromised by the resulting ordering rule violations.

For the use model where a single Requester uses both Translated and Untranslated addresses with peer-to-peer Requests, again it is recommended that software not enable ACS Direct Translated P2P unless software is certain that proper operation is not compromised by the resulting ordering rule violations. This requires a detailed analysis of the peer-to-peer communications models being used, and is beyond the scope of this specification.

---

*Add to Chapter 7*

## 7.10.2. Uncorrectable Error Status Register (Offset 04h)

...

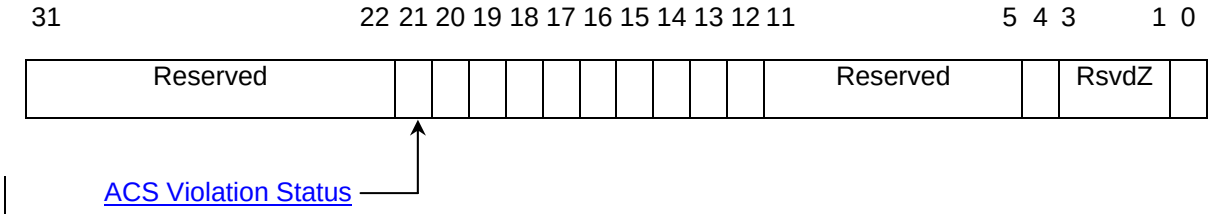


**Figure 7-28 Uncorrectable Error Status Register**

**Table 7-24 Uncorrectable Error Status Register**

| Bit Location | Register Description | Attributes | Default Value |
|---|---|---|---|
| … | **...** | … | … |
| 21 | **ACS Violation Status** | RW1CS | 0 |

## 7.10.3. Uncorrectable Error Mask Register (Offset 08h)

...

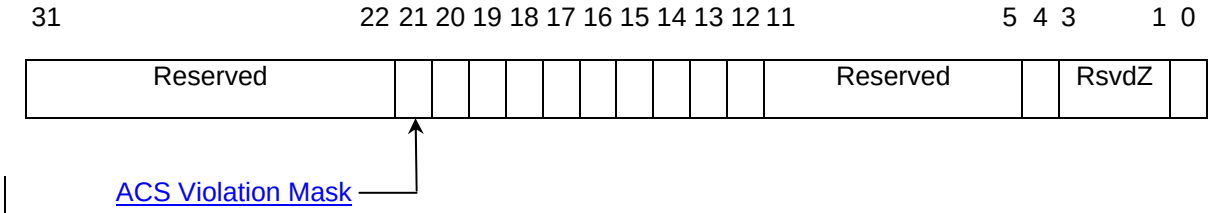


**Figure 7-29 Uncorrectable Error Mask Register**

**Table 7-25 Uncorrectable Error Mask Register**

| Bit Location | Register Description | Attributes | Default Value |
|---|---|---|---|
| … | **...** | … | … |
| 21 | **ACS Violation Mask** | RWS | 0 |

## 7.10.4. Uncorrectable Error Severity Register (Offset 0Ch)

...

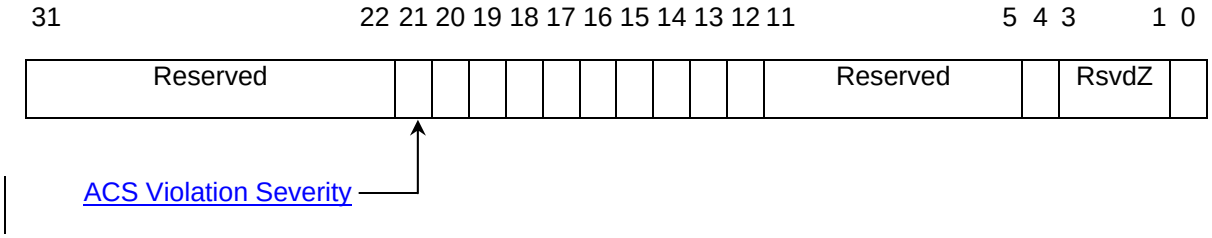


**Figure 7-30 Uncorrectable Severity Mask Register**

**Table 7-26 Uncorrectable Severity Mask Register**

| Bit Location | Register Description | Attributes | Default Value |
|---|---|---|---|
| … | **...** | … | … |
| 21 | **ACS Violation Severity** | RWS | 0 |

# 7.16. ACS Extended Capability

The ACS Extended Capability is an optional capability that provides enhanced access controls. See Section 6.11. This capability may be implemented by a Root Port , a Switch Downstream Port, or a multi-function device Function. It is never applicable to a PCI Express to PCI Bridge or Root Complex Event Collector. It is not applicable to a Switch Upstream Port unless that Switch Upstream Port is a Function in a multi-function device.

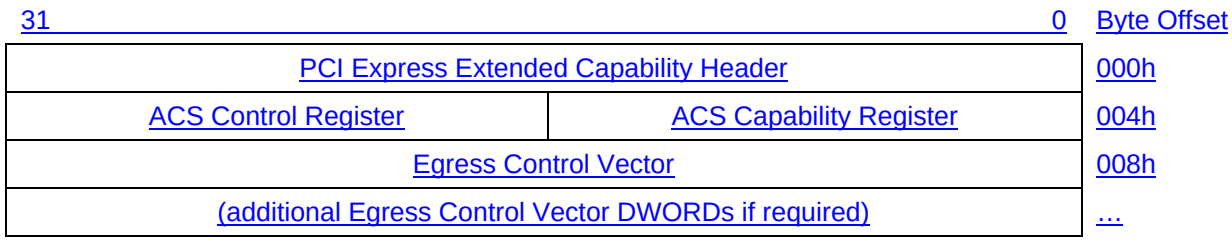


**Figure 7-xx ACS Extended Capability**

## 7.16.1. ACS Extended Capability Header (Offset 00h)

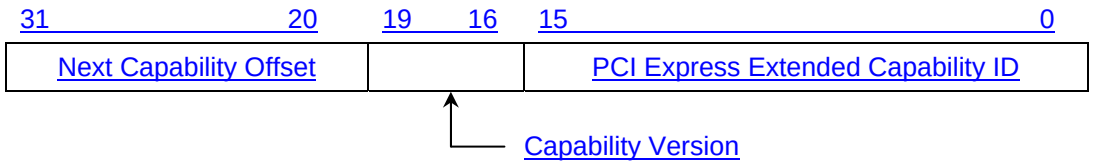


**Figure 7-xx ACS Extended Capability Header**

**Table 7-xx ACS Extended Capability Header**

| Bit Location | Register Description | Attributes |
|---|---|---|
| 15:0 | **PCI Express Extended Capability ID** – This field is a PCI-SIG defined ID number that indicates the nature and format of the extended capability.<br><br>PCI Express Extended Capability ID for the ACS Extended Capability is 000Dh. | RO |
| 19:16 | **Capability Version** – This field is a PCI-SIG defined version number that indicates the version of the capability structure present.<br><br>Must be 1h for this version of the specification. | RO |
| 31:20 | **Next Capability Offset** – This field contains the offset to the next PCI Express Extended Capability structure or 000h if no other items exist in the linked list of capabilities. | RO |

## 7.16.2. ACS Capability Register (Offset 04h)

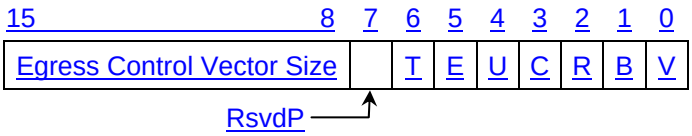


**Figure 7-xx ACS Capability Register**

**Table 7-xx ACS Capability Register**

| Bit Location | Register Description | Attributes |
|---|---|---|
| 0 | **ACS Source Validation (V)** – Required for Root Ports and Switch Downstream Ports; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS Source Validation. | RO |
| 1 | **ACS Translation Blocking (B)** – Required for Root Ports and Switch Downstream Ports; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS Translation Blocking. | RO |
| 2 | **ACS P2P Request Redirect (R)** – Required for Root Ports that support peer-to-peer traffic with other Root Ports; required for Switch Downstream Ports; required for multi-function device Functions that support peer-to-peer traffic with other Functions; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS P2P Request Redirect. | RO |
| 3 | **ACS P2P Completion Redirect (C)** – Required for all Functions that support ACS P2P Request Redirect; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS P2P Completion Redirect. | RO |
| 4 | **ACS Upstream Forwarding (U) –** Required for Root Ports if the RC supports Redirected Request Validation; required for Switch Downstream Ports; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS Upstream Forwarding. | RO |
| 5 | **ACS P2P Egress Control (E) –** Optional for Root Ports, Switch Downstream Ports, and multi-function device Functions; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS P2P Egress Control. | RO |
| 6 | **ACS Direct Translated P2P (T)** – Required for Root Ports that support Address Translation Services (ATS) and also support peer-to-peer traffic with other Root Ports; required for Switch Downstream Ports; must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS Direct Translated P2P. | RO |
| 15:8 | **Egress Control Vector Size** – Encodings 01h-FFh directly indicate the number of applicable bits in the Egress Control Vector; the encoding 00h indicates 256 bits.<br><br>If the ACS P2P Egress Control (E) bit in this register is 0b, the value of the size field is undefined, and the vector is not required to be present. | HwInit |

## 7.16.3. ACS Control Register (Offset 06h)

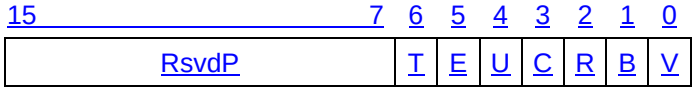


**Figure 7-xx ACS Control Register**

**Table 7-xx ACS Control Register**

| Bit Location | Register Description | Attributes |
|---|---|---|
| 0 | **ACS Source Validation Enable (V)** – When set, the component validates the Bus Number from the Requester ID of upstream Requests against the secondary / subordinate Bus Numbers.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS Source Validation functionality is not implemented. | RW |
| 1 | **ACS Translation Blocking Enable (B)** – When set, the component blocks all upstream Memory Requests whose Address Translation (AT) field is not set to the default value.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS Translation Blocking functionality is not implemented. | RW |
| 2 | **ACS P2P Request Redirect Enable (R)** – In conjunction with ACS P2P Egress Control and ACS Direct Translated P2P mechanisms, determines when the component redirects peer-to-peer Requests upstream. See Section 6.11.3. Note that with Downstream Ports, the R bit only applies to upstream Requests arriving at the Downstream Port, and whose normal routing targets a different Downstream Port.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS P2P Request Redirect functionality is not implemented. | RW |
| 3 | **ACS P2P Completion Redirect Enable (C)** – Determines when the component redirects peer-to-peer Completions upstream; applicable only to Read Completions whose Relaxed Ordering Attribute is clear.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS P2P Completion Redirect functionality is not implemented. | RW |
| 4 | **ACS Upstream Forwarding Enable (U)** – When set, the component forwards upstream any Request or Completion TLPs it receives that were redirected upstream by a component lower in the hierarchy. Note that the U bit only applies to upstream TLPs arriving at a Downstream Port, and whose normal routing targets the same Downstream Port.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS Upstream Forwarding functionality is not implemented. | RW |
| 5 | **ACS P2P Egress Control Enable (E) –** In conjunction with the Egress Control Vector plus the ACS P2P Request Redirect and ACS Direct Translated P2P mechanisms, determines when to allow, disallow, or redirect peer-to-peer Requests. See Section 6.11.3.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS P2P Egress Control functionality is not implemented. | RW |

| | | |
|---|---|---|
| 6 | **ACS Direct Translated P2P Enable (T) –** When set, overrides the ACS P2P Request Redirect and ACS P2P Egress Control mechanisms with peer-to-peer Memory Requests whose Address Translation (AT) field indicates a Translated address. See Section 6.11.3.<br><br>Default value of this field is 0b. Must be hardwired to 0b if the ACS Direct Translated P2P functionality is not implemented. | RW |

## 7.16.4. Egress Control Vector (Offset 08h)

The Egress Control Vector is a read-write register that contains a bit-array. The number of bits in the register is specified by the Egress Control Vector Size field, and the register spans multiple DWORDs if required. If the ACS P2P Egress Control bit in the ACS Capability register is 0b, the Egress Control Vector Size field is undefined, and the Egress Control Vector is not required to be present.

For the general case of an Egress Control Vector spanning multiple DWORDs, the DWORD offset and bit number within that DWORD for a given arbitrary bit *K* are specified by the formulas:

DWORD offset = 08h + (K div[12] 32) * 4

DWORD bit# = K mod[13] 32

Bits in a DWORD beyond those specified by the Egress Control Vector Size field are RsvdP.

For Root Ports and Switch Downstream Ports, each bit in the bit-array always corresponds to a Port Number. Otherwise, for Functions[14] within a multi-function device, each bit in the bit-array corresponds to a Function Number. For example, access to Function 2 is controlled by bit number 2 in the bit-array. For both Port Number cases and Function Number cases, the bit corresponding to the Function that implements this Extended Capability structure must be hardwired to 0b.

With RCs, some Port Numbers may refer to internal Ports instead of Root Ports. For Root Ports in such RCs, each bit in the bit-array that corresponds to an internal Port must be hardwired to 0b.

[12] Div is an integer divide with truncation.

[13] Mod is the remainder from an integer divide.

[14] Including Switch Upstream Ports.

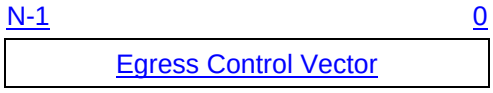


**Figure 7-xx Egress Control Vector**

**Table 7-xx Egress Control Vector**

| Bit Location | Register Description | Attributes |
|---|---|---|
| N-1:0 | **Egress Control Vector –** A N-bit bit-array configured by software. When a given bit is set, peer-to-peer Requests targeting the associated Port or Function are blocked or redirected (if enabled). See Section 6.11.3.<br><br>Default value of this field is 0. | RW |

The following examples illustrate how the vector might be configured:

- For an 8-Port Switch, each Port will have a separate vector indicating which Downstream Egress Ports it may forward Requests to.
  - Port 1 being not allowed to communicate with any other Downstream Ports would be configured as: 1111 1100b with 0b indicating in bit 0 corresponds to the upstream Port and a 0b in bit 1 represents the Ingress Port hardwired to 0b as well.
  - Port 2 being allowed to communicate with Ports 3, 5, and 7 would be configured as: 0101 0010b.
- For a 4-Function device, each Function will have a separate vector that indicates which Function it may forward Requests to.
  - Function 0 being not allowed to communicate with any other Functions would be configured as: 1110b with 0b in bit 0 corresponding to Function 0.
  - Function 1 being allowed to communicate with Functions 2 and 3 would be configured as: 0001b with a 0b in bit 1 corresponding to Function 1 hardwired to 0b as well.