

NVM Express Technical Errata

Errata ID	017
Change Date	7/13/2011
Affected Spec Ver.	NVM Express 1.0b
Corrected Spec Ver.	

Submission info

Name	Company	Date
Kwok Kong	IDT	6/30/2011

This erratum clarifies the information returned in Dword 0 of the command completion for each Feature Identifier using the Get Features command.

Add a second paragraph to section 5.12.1.1 as shown below:

This Feature controls command arbitration. Refer to section 4.7 for command arbitration details. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 74 are returned in Dword 0 of the completion queue entry for that command.

Add a second paragraph to section 5.12.1.2 as shown below:

This Feature allows the host to configure the power state. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 75 are returned in Dword 0 of the completion queue entry for that command.

Add a new fourth paragraph to section 5.12.1.3 as shown below:

This feature indicates the type and attributes of LBA ranges that are part of the specified namespace. The LBA range information may be used by a driver to determine if it may utilize a particular LBA range; the information is not exposed to higher level software.

This is optional information that is not required for proper behavior of the system. However, it may be utilized to avoid unintended software issues. For example, if the LBA range indicates that it is a RAID volume then a driver that does not have RAID functionality should not utilize that LBA range (including not overwriting the LBA range). The optional information may be utilized by the driver to determine whether the LBA Range should be exposed to higher level software.

The LBA Range Type uses Command Dword 11 and specifies the type and attribute information in the data structure indicated in Figure 77. The data structure is 4096 bytes in size and shall be physically contiguous.

If a Get Features command is issued for this Feature, the attributes specified in Figure 76 are returned in Dword 0 of the completion queue entry and the LBA Range Type data structure specified in Figure 77 is returned in the data buffer for that command.

Add a second paragraph to section 5.12.1.4 as shown below:

This Feature indicates the threshold for the temperature of the overall device (controller and NVM included) in units of Kelvin. If this temperature is exceeded, then an asynchronous event may be issued to the host. The host should configure this feature prior to enabling asynchronous event notification for the temperature exceeding the threshold. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 78 are returned in Dword 0 of the completion queue entry for that command.

Add a second paragraph to section 5.12.1.5 as shown below:

This Feature indicates the threshold for the temperature of the overall device (controller and NVM included) in units of Kelvin. If this temperature is exceeded, then an asynchronous event may be issued to the host. The host should configure this feature prior to enabling asynchronous event notification for the temperature exceeding the threshold. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 79 are returned in Dword 0 of the completion queue entry for that command.

Add a third paragraph to section 5.12.1.6 as shown below:

This Feature controls the volatile write cache, if present, on the controller. If a volatile write cache is supported, then this feature shall be supported. The attributes are indicated in Command Dword 11.

Note: If the controller is able to guarantee that data present in a write cache is written to non-volatile media on loss of power, then that write cache is considered non-volatile and this setting does not apply to that write cache. In that case, this setting has no effect.

If a Get Features command is issued for this Feature, the attributes specified in Figure 80 are returned in Dword 0 of the completion queue entry for that command.

Add a second paragraph to section 5.12.1.7 as shown below:

This Feature indicates the number of queues that the host requests for this controller. This feature shall only be issued during initialization prior to creation of any I/O Submission and/or Completion Queues. The value allocated shall not change between resets. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 82 are returned in Dword 0 of the completion queue entry for that command.

Add a third paragraph to section 5.12.1.8 as shown below:

This Feature configures interrupt coalescing settings. This Feature applies only to the I/O Queues. It is recommended that interrupts for commands that complete in error are not coalesced. The settings are specified in Command Dword 11.

The controller may delay an interrupt if it detects that interrupts are already being processed for this vector. Specifically, if the Completion Queue Head Doorbell register is being updated that is associated with a particular interrupt vector, then the controller has a positive indication that completion queue entries are already being processed. In this case, the aggregation time and/or the aggregation threshold may be reset/restarted upon the associated register write. This may result in interrupts being delayed indefinitely in certain workloads where the aggregation time or aggregation threshold are non-zero.

If a Get Features command is issued for this Feature, the attributes specified in Figure 83 are returned in Dword 0 of the completion queue entry for that command.

Add a third paragraph to section 5.12.1.9 as shown below:

This Feature configures settings specific to a particular interrupt vector. The settings are specified in Command Dword 11.

By default, coalescing settings are enabled for each interrupt vector. Interrupt coalescing is not supported for the Admin Completion Queue.

If a Get Features command is issued for this Feature, the attributes specified in Figure 84 are returned in Dword 0 of the completion queue entry for that command.

Add a second paragraph to section 5.12.1.10 as shown below:

This Feature controls write atomicity. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 85 are returned in Dword 0 of the completion queue entry for that command.

Add a second paragraph to section 5.12.1.11 as shown below:

This Feature controls the events that trigger an asynchronous event notification to the host. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 86 are returned in Dword 0 of the completion queue entry for that command.

Add a second paragraph to section 5.12.1.12 as shown below:

This Feature is a software progress marker. The software progress marker is persistent across power states. For additional details, refer to section 7.6.1.1. This information may be used to indicate to an OS software driver whether there have been issues with the OS successfully loading. The attributes are indicated in Command Dword 11.

If a Get Features command is issued for this Feature, the attributes specified in Figure 87 are returned in Dword 0 of the completion queue entry for that command.

Disposition log

6/30/2011	Erratum captured.
7/13/2011	Editorial updates.
8/30/2011	Erratum ratified.

Technical input submitted to the NVMHCI Workgroup is subject to the terms of the NVMHCI Contributor's agreement.