

ATA-7 Write DMA Queued FUA Ext Problem

To: T13 Technical committee
 From: John Masiewicz
 Western Digital
 20511 Lake Forest Drive
 Lake Forest, CA 92639
 Phone: 949-672-7686
 Email: john.masiewicz@wdc.com
 Date: June 19, 2003

1 Introduction

Western Digital believes there is a problem with the requirement under the "description" paragraph for the Write DMA Queued FUA EXT command (d1532v1r2a, Volume 1 clause 6.62.8, and also mentioned in clause 4.20).

Clause 4.20 (last paragraph) states:

“The WRITE DMA QUEUED FUA EXT command is unique in that it shall not be released. When this command is received, the user data shall be written to the device media before ending status for the command is reported regardless of the state of any write cache or queue. A queue shall not be aborted”.

Clause 6.62.8 states:

“The WRITE DMA QUEUED FUA EXT command provides the same function as the WRITE DMA EXT command. It is an Overlapped feature set command and when issued it shall not cause an existing queue to be aborted. However, this command shall not be released and regardless of whether write caching in the device is enabled or not, the user data shall be written to the media before ending status for the command is reported.”

There is a fundamental problem if the drive is not allowed to release on this command. Take the following command sequence:

```

RRR          Queued Read #1 - released
  WWW       Queued Write #1 - released
    RRR     Queued Read #2 - released
WWWWWWWWWW Write DMA Queued FUA EXT
  
```

If the drive released (did not transfer any data) on the first three commands, it's in real trouble if it can't release on the fourth command. It can't provide the host with the proper data for the two released Queued reads. Furthermore, after writing the data for the fourth command to the media, it would have to request Service for the first write, take in the data from the host, and then throw it away. One embodiment of a Queuing implementation would be for the device to release on the fourth command, get the first three commands serviced, then request Service for the fourth command, get the data from the host, and hold BSY until the data is written to the media.

The description should reflect the intention of the command that the device must write the user data to the media before ending status for the command is reported. Stated another way, WHEN the drive does decide to take the data from the host, it must hold BSY until the data is successfully written on the media. It cannot just take the data into its buffer and then clear BSY.

2 Proposed Specification Changes

There is no clear reasoning behind the restriction to not release a FUA command. It is an unnecessary restriction and causes the severe concurrency problems noted above. We recommend the simplest change, which is to remove this restriction, and allow the device to release the command.

Change Clause 4.20 (last paragraph) to:

~~“The WRITE DMA QUEUED FUA EXT command is unique in that it shall not be released. When this command is received, the user data shall be written to the device media before ending status for the command is reported regardless of the state of any write cache or queue. A queue shall not be aborted”.~~

Change Clause 6.62.8 to:

“The WRITE DMA QUEUED FUA EXT command provides the same function as the WRITE DMA EXT command. It is an Overlapped feature set command and when issued it shall not cause an existing queue to be aborted. However, ~~this command shall not be released and regardless of whether write caching in the device is enabled or not, the user data shall be written to the media before ending status for the command is reported.~~”

Revision 1 changes:

The following Changes are added in Revision 1 of this document. The original Write DMA QUEUED FUA EXT was copied from WRITE DMA EXT, and should have been copied from WRITE DMA QUEUED EXT. The above change and agreement was that the FUA command could be released, which requires the command to follow the Queued protocol. The following sections show the deleted material with strikeouts and the corrected material following.

The replaced section is from d1532v1r3 section 6.62.

~~6.62 WRITE DMA QUEUED FUA EXT~~

~~6.62.1 Command code~~

~~3Eh~~

~~6.62.2 Feature set~~

~~Overlapped feature set and 48-bit Address feature set~~

~~?? Optional for devices implementing the 48-bit Address feature set.~~

~~?? Use prohibited for devices implementing the PACKET Command feature set.~~

~~6.62.3 Protocol~~

~~DMA QUEUED(see clause [Error! Reference source not found.](#)).~~

6.62.4 Inputs

Register		7	6	5	4	3	2	1	0
Features	Current Previous	Reserved Reserved							
Sector Count	Current Previous	Sector count (7:0) Sector count (15:8)							
LBA Low	Current Previous	LBA (7:0) LBA (31:24)							
LBA Mid	Current Previous	LBA (15:8) LBA (39:32)							
LBA High	Current Previous	LBA (23:16) LBA (47:40)							
Device		obs	LBA	obs	DEV	Reserved			
Command		3Eh							
NOTE ? The value indicated as Current is the value most recently written to the register. The value indicated as Previous is the value that was in the register before the most recent write to the register.									

Sector Count Current -

number of sectors to be transferred low order, bits (7:0).

Sector Count Previous -

number of sectors to be transferred high order, bits (15:8). 0000h in the Sector Count register specifies that 65,536 sectors are to be transferred.

LBA Low Current -

LBA (7:0).

LBA Low Previous -

LBA (31:24).

LBA Mid Current -

LBA (15:8).

LBA Mid Previous -

LBA (39:32).

LBA High Current -

LBA (23:16).

LBA High Previous -

LBA (47:40).

Device -

the LBA bit shall be set to one to specify the address is an LBA.
DEV shall specify the selected device.

6.62.5 Normal outputs

Register		7	6	5	4	3	2	1	0
Error		na							
Sector Count	HOB = 0				Reserved				
	HOB = 1				Reserved				
LBA Low	HOB = 0				Reserved				
	HOB = 1				Reserved				
LBA Mid	HOB = 0				Reserved				
	HOB = 1				Reserved				
LBA High	HOB = 0				Reserved				
	HOB = 1				Reserved				
Device		obs	na	obs	DEV	Reserved			
Status		BSY	DRDY	DF	na	DRQ	na	na	ERR
NOTE ? HOB = 0 indicates the value read by the host when the HOB bit of the Device Control register is cleared to zero. HOB = 1 Indicates the value read by the host when the HOB bit of the Device Control register is set to one.									

Device register –

DEV shall indicate the selected device.

Status register –

BSY shall be cleared to zero indicating command completion.

DRDY shall be set to one.

DF (Device Fault) shall be cleared to zero.

DRQ shall be cleared to zero.

ERR shall be cleared to zero.

6.62.6 Error outputs

An unrecoverable error encountered during the execution of this command results in the termination of the command. The Command Block registers contain the address of the sector where the first unrecoverable error occurred. The amount of data transferred is indeterminate.

Register		7	6	5	4	3	2	1	0
Error		ICRC	WP	MC	IDNF	MCR	ABRT	NM	obs
Sector Count	HOB = 0	Reserved							
	HOB = 1	Reserved							
LBA Low	HOB = 0				LBA (7:0)				
	HOB = 1				LBA (31:24)				
LBA Mid	HOB = 0				LBA (15:8)				
	HOB = 1				LBA (39:32)				
LBA High	HOB = 0				LBA (23:16)				
	HOB = 1				LBA (47:40)				
Device		obs	na	obs	DEV	Reserved			
Status		BSY	DRDY	DF	na	DRQ	na	na	ERR
NOTE ? HOB = 0 indicates the value read by the host when the HOB bit of the Device Control register is cleared to zero. HOB = 1 Indicates the value read by the host when the HOB bit of the Device Control register is set to one.									

Error register-

~~ICRC shall be set to one if an interface CRC error has occurred during an Ultra-DMA data transfer. The content of this bit is not applicable for Multiword-DMA transfers.~~

~~WP shall be set to one if the media in a removable media device is write-protected.~~

~~MC shall be set to one if the media in a removable media device changed since the issuance of the last command. The device shall clear the device internal media change detected state.~~

~~IDNF shall be set to one if a user-accessible address could not be found IDNF shall be set to one if an address outside of the range of user-accessible addresses is requested if command aborted is not returned.~~

~~MCR shall be set to one if a media change request has been detected by a removable media device. This bit is only cleared by a GET MEDIA STATUS or a media access command.~~

~~ABRT shall be set to one if this command is not supported or if an error, including an ICRC error, has occurred during an Ultra-DMA data transfer. ABRT may be set to one if the device is not able to complete the action requested by the command. ABRT shall be set to one if an address outside of the range of user-accessible addresses is requested if IDNF is not set to one.~~

~~NM shall be set to one if no media is present in a removable media device.~~

LBA Low-

~~LBA (7:0) of the address of the first unrecoverable error when read with Device Control register HOB bit cleared to zero.~~

~~LBA (31:24) of the address of the first unrecoverable error when read with Device Control register HOB bit set to one.~~

LBA Mid-

~~LBA (15:8) of the address of the first unrecoverable error when read with Device Control register HOB bit cleared to zero.~~

~~LBA (39:32) of the address of the first unrecoverable error when read with Device Control register HOB bit set to one.~~

LBA High-

~~LBA (23:16) of the address of the first unrecoverable error when read with Device Control register HOB bit cleared to zero.~~

~~LBA (47:40) of the address of the first unrecoverable error when read with Device Control register HOB bit set to one.~~

Device register-

~~DEV shall indicate the selected device.~~

Status register-

~~BSY shall be cleared to zero indicating command completion.~~

~~DRDY shall be set to one.~~

~~DF (Device Fault) shall be set to one if a device fault has occurred.~~

~~DRQ shall be cleared to zero.~~

~~ERR shall be set to one if an Error register bit is set to one.~~

6.62.7 Prerequisites

~~DRDY set to one. The host shall initialize the DMA channel.~~

6.62.8 Description

~~The WRITE DMA QUEUED FUA EXT command provides the same function as the WRITE DMA EXT command. It is an Overlapped feature set command and when issued it shall not cause an existing queue to be aborted. However, this command shall not be released and regardless of whether write caching in the device is enabled or not, the user data shall be written to the media before ending status for the command is reported.~~

6.62 WRITE DMA QUEUED FUA EXT

6.62.1 Command code

3Eh

6.62.2 Feature set

Overlapped feature set and 48-bit Address feature set

?? Mandatory for devices implementing the Overlapped feature set and the 48-bit Address feature set and not implementing the PACKET Command feature set.

?? Use prohibited for devices implementing the PACKET Command feature set.

6.62.3 Protocol

DMA QUEUED (see clause **Error! Reference source not found.**).

6.62.4 Inputs

Register		7	6	5	4	3	2	1	0	
Features	Current Previous	Sector count (7:0) Sector count (15:8)								
Sector Count	Current	Tag					Reserved			
	Previous	Reserved								
LBA Low	Current Previous	LBA (7:0) LBA (31:24)								
LBA Mid	Current Previous	LBA (15:8) LBA (39:32)								
LBA High	Current Previous	LBA (23:16) LBA (47:40)								
Device		obs	LBA	obs	DEV	Reserved				
Command		3Eh								
NOTE ? The value indicated as Current is the value most recently written to the register. The value indicated as Previous is the value that was in the register before the most recent write to the register.										

Features Current -
number of sectors to be transferred low order, bits (7:0).

Features Previous -
number of sectors to be transferred high order, bits (15:8). 0000h in the Features register specifies that 65,536 sectors are to be transferred.

Sector Count Current -
if the device supports command queuing, bits (7:3) contain the Tag for the command being delivered. A Tag value may be any value between 0 and 31 regardless of the queue depth supported. If queuing is not supported, this register shall be set to the value 00h.

Sector Count Previous -
Reserved

LBA Low Current -

- LBA (7:0).
- LBA Low Previous - LBA (31:24).
- LBA Mid Current - LBA (15:8).
- LBA Mid Previous - LBA (39:32).
- LBA High Current - LBA (23:16).
- LBA High Previous - LBA (47:40).

Device -

the LBA bit shall be set to one to specify the address is an LBA.
DEV shall specify the selected device.

6.62.5 Normal outputs

6.62.5.1 Data transmission

Data transfer may occur after receipt of the command or may occur after the receipt of a SERVICE command. When the device is ready to transfer data requested by a data transfer command, the device sets the following register content to initiate the data transfer.

Register		7	6	5	4	3	2	1	0	
Error		na								
Sector Count	HOB = 0	Tag					REL	I/O	C/D	
	HOB = 1	Reserved								
LBA Low	HOB = 0	Reserved								
	HOB = 1	Reserved								
LBA Mid	HOB = 0	Reserved								
	HOB = 1	Reserved								
LBA High	HOB = 0	Reserved								
	HOB = 1	Reserved								
Device		obs	na	obs	DEV	Reserved				
Status		BSY	DRDY	DF	na	DRQ	na	na	ERR	
NOTE ? HOB = 0 indicates the value read by the host when the HOB bit of the Device Control register is cleared to zero. HOB = 1 Indicates the value read by the host when the HOB bit of the Device Control register is set to one.										

Sector Count (when HOB of the Device Control register is cleared to zero) -

Tag - This field contains the command Tag for the command. A Tag value may be any value between 0 and 31 regardless of the queue depth supported. If the device does not support command queuing or overlap is disabled, this register shall be set to the value 00h.

REL - Shall be cleared to zero.

I/O - Shall be cleared to zero indicating the transfer is from the host.

C/D - Shall be cleared to zero indicating the transfer of data.

Device register -

DEV shall indicate the selected device.

Status register -

BSY shall be cleared to zero.

DRDY shall be set to one.

DF (Device Fault) shall be cleared to zero.

DRQ shall be cleared to zero.

ERR shall be cleared to zero.

6.62.5.2 Bus release

If the device performs a bus release before transferring data for this command, the register content upon performing a bus release shall be as shown below.

Register		7	6	5	4	3	2	1	0
Error		na							
Sector Count	HOB = 0	Tag				REL	I/O	C/D	
	HOB = 1	Reserved							
LBA Low	HOB = 0	Reserved							
	HOB = 1	Reserved							
LBA Mid	HOB = 0	Reserved							
	HOB = 1	Reserved							
LBA High	HOB = 0	Reserved							
	HOB = 1	Reserved							
Device		obs	na	obs	DEV	Reserved			
Status		BSY	DRDY	DF	SERV	DRQ	na	na	ERR
NOTE ? HOB = 0 indicates the value read by the host when the HOB bit of the Device Control register is cleared to zero. HOB = 1 Indicates the value read by the host when the HOB bit of the Device Control register is set to one.									

Sector Count (when HOB of the Device Control register is cleared to zero) -

Tag -This field contains the command Tag for the command. A Tag value may be any value between 0 and 31 regardless of the queue depth supported. If the device does not support command queuing or overlap is disabled, this register shall be set to the value 00h.

REL - Shall be set to one.

I/O - Shall be cleared to zero.

C/D - Shall be cleared to zero indicating the transfer of data.

Device register –

DEV shall indicate the selected device.

Status register –

BSY shall be cleared to zero.

DRDY shall be set to one.

DF (Device Fault) shall be cleared to zero.

SERV (Service) shall be cleared to zero when no other queued command is ready for service.

SERV shall be set to one when another queued command is ready for service. SERV shall be set to one when the device has prepared this command for service.

DRQ shall be cleared to zero.

ERR shall be cleared to zero.

6.62.5.3 Service request

When the device is ready to transfer data or complete a command after the command has performed a bus release, the device shall set the SERV bit to one and not change the state of any other register bit (see clause **Error! Reference source not found.**). When the SERVICE command is received, the device shall set outputs as described in data transfer, command completion, or error outputs depending on the service the device requires.

6.62.5.4 Command completion

When the transfer of all requested data has occurred without error, the register content shall be as shown below.

Register		7	6	5	4	3	2	1	0
Error		na							
Sector Count	HOB = 0	Tag				REL	I/O	C/D	
	HOB = 1	Reserved							
LBA Low	HOB = 0	Reserved							
	HOB = 1	Reserved							
LBA Mid	HOB = 0	Reserved							
	HOB = 1	Reserved							
LBA High	HOB = 0	Reserved							
	HOB = 1	Reserved							
Device		obs	na	obs	DEV	Reserved			
Status		BSY	DRDY	DF	SERV	DRQ	na	na	ERR
NOTE ? HOB = 0 indicates the value read by the host when the HOB bit of the Device Control register is cleared to zero. HOB = 1 Indicates the value read by the host when the HOB bit of the Device Control register is set to one.									

Sector Count (when HOB of the Device Control register is cleared to zero) -

Tag -This field contains the command Tag for the command. A Tag value may be any value between 0 and 31 regardless of the queue depth supported. If the device does not support command queuing or overlap is disabled, this register shall be set to the value 00h.

REL - Shall be cleared to zero.

I/O - Shall be set to one.

C/D - Shall be set to one.

Device register –

DEV shall indicate the selected device.

Status register –

BSY shall be cleared to zero indicating command completion.

DRDY shall be set to one.

DF (Device Fault) shall be cleared to zero.

SERV (Service) shall be cleared to zero when no other queued command is ready for service.

SERV shall be set to one when another queued command is ready for service.

DRQ shall be cleared to zero.

ERR shall be cleared to zero.

6.62.6 Error outputs

The Sector Count register contains the Tag for this command if the device supports command queuing. The device shall return command aborted if the command is not supported or if the device has not had overlapped interrupt enabled. The device shall return command aborted if the device supports command queuing and the Tag is invalid. An unrecoverable error encountered during the execution of this command results in the termination of the command and the Command Block registers contain the sector where the first unrecoverable error occurred. If a queue existed, the unrecoverable error shall cause the queue to abort.

Register		7	6	5	4	3	2	1	0
Error		ICRC	WP	MC	IDNF	MCR	ABRT	NM	obs
Sector Count	HOB = 0	Tag					REL	I/O	C/D
	HOB = 1	Reserved							
LBA Low	HOB = 0	LBA (7:0)							
	HOB = 1	LBA (31:24)							
LBA Mid	HOB = 0	LBA (15:8)							
	HOB = 1	LBA (39:32)							
LBA High	HOB = 0	LBA (23:16)							
	HOB = 1	LBA (47:40)							
Device		obs	na	obs	DEV	Reserved			
Status		BSY	DRDY	DF	SERV	DRQ	na	na	ERR
NOTE ? HOB = 0 indicates the value read by the host when the HOB bit of the Device Control register is cleared to zero. HOB = 1 Indicates the value read by the host when the HOB bit of the Device Control register is set to one.									

Error register -

ICRC shall be set to one if an interface CRC error has occurred during an Ultra DMA data transfer. The content of this bit is not applicable for Multiword DMA transfers.

WP shall be set to one if the media in a removable media device is write protected.

MC shall be set to one if the media in a removable media device changed since the issuance of the last command. The device shall clear the device internal media change detected state.

IDNF shall be set to one if a user-accessible address could not be found. IDNF shall be set to one if an address outside of the range of user-accessible addresses is requested if command aborted is not returned.

MCR shall be set to one if a media change request has been detected by a removable media device. This bit is only cleared by a GET MEDIA STATUS or a media access command.

ABRT shall be set to one if this command is not supported or if an error, including an ICRC error, has occurred during an Ultra DMA data transfer. ABRT may be set to one if the device is not able to complete the action requested by the command. ABRT shall be set to one if an address outside of the range of user-accessible addresses is requested if IDNF is not set to one.

NM shall be set to one if no media is present in a removable media device.

Sector Count (when HOB of the Device Control register is cleared to zero) -

Tag - This field contains the command Tag for the command. A Tag value may be any value between 0 and 31 regardless of the queue depth supported. If the device does not support command queuing or overlap is disabled, this register shall be set to the value 00h.

REL - Shall be cleared to zero.

I/O - Shall be set to one.

C/D - Shall be set to one.

LBA Low -

LBA (7:0) of the address of the first unrecoverable error when read with Device Control register HOB bit cleared to zero.

LBA (31:24) of the address of the first unrecoverable error when read with Device Control register HOB bit set to one.

LBA Mid -

LBA (15:8) of the address of the first unrecoverable error when read with Device Control register HOB bit cleared to zero.

LBA (39:32) of the address of the first unrecoverable error when read with Device Control register HOB bit set to one.

LBA High -

LBA (23:16) of the address of the first unrecoverable error when read with Device Control register HOB bit cleared to zero.

LBA (47:40) of the address of the first unrecoverable error when read with Device Control register HOB bit set to one.

Device register -

DEV shall indicate the selected device.

Status register -

BSY shall be cleared to zero indicating command completion.

DRDY shall be set to one.

DF (Device Fault) shall be set to one if a device fault has occurred.

DRQ shall be cleared to zero.

ERR shall be set to one if an Error register bit is set to one.

6.62.7 Prerequisites

DRDY set to one. The host shall initialize the DMA channel.

6.62.8 Description

This command executes in a similar manner to a WRITE DMA EXT command. The device may perform a bus release the bus or may execute the data transfer without performing a bus release if the data is ready to transfer.

If the device performs a bus release, the host shall reselect the device using the SERVICE command.

Once the data transfer is begun, the device shall not perform a bus release until the entire data transfer has been completed.

The WRITE DMA QUEUED FUA EXT command provides the same function as the WRITE DMA EXT command. It is an Overlapped feature set command and when issued it shall not cause an existing queue to be aborted. However, regardless of whether write caching in the device is enabled or not, the user data shall be written to the media before ending status for the command is reported.