SFF-8609

Specification for

Management Interface for Drive Conditions

Rev 1.0 July 07, 2017

Secretariat: SFF TA TWG

Abstract: This specification defines a method to monitor various storage device conditions by embedding storage device management data in serial packets that are transmitted across the storage device ready/activity signal.

Server and storage products that contain multiple HDDs (Hard Disk Drives) and SSDs (Solid State Disks) present a challenge for the effective monitoring of various storage device conditions such as thermal monitoring or power utilization. In-band commands such as the SCSI 'Log Sense' command have the potential to interrupt data flow, and in multiple node systems the node which owns the data path is not necessarily the node responsible for chassis management.

This specification provides a common reference for systems manufacturers, system integrators, and suppliers implementing a device management protocol using the storage device ready/activity signal.

This specification is made available for public review, and written comments are solicited from readers. Comments received by the members will be considered for inclusion in future revisions of this specification.

POINTS OF CONTACT:

Bill Lynn
Dell Technologies
One Dell Way #RR5-30
Round Rock TX 78682
Ph: 512-723-5528
Email: William.Lynn@Dell.com

Chairman SFF TA TWG
Email: SFF-Chair@snia.org

Foreword

The development work on this specification was done by the SNIA SFF TWG, an industry group. Since its formation as the SFF Committee in August 1990, the membership has included a mix of companies which are leaders across the industry.

When 2 1/2" diameter disk drives were introduced, there was no commonality on external dimensions e.g. physical size, mounting locations, connector type, connector location, between vendors. The SFF Committee provided a forum for system integrators and vendors to define the form factor of disk drives.

During their definition, other activities were suggested because participants in SFF faced more challenges than the form factors. In November 1992, the charter was expanded to address any issues of general interest and concern to the storage industry. The SFF Committee became a forum for resolving industry issues that are either not addressed by the standards process or need an immediate solution.

In July 2016, the SFF Committee transitioned to SNIA (Storage Networking Industry Association), as a TA (Technology Affiliate) TWG (Technical Work Group).

Industry consensus is not a requirement to publish a specification because it is recognized that in an emerging product area, there is room for more than one approach. By making the documentation on competing proposals available, an integrator can examine the alternatives available and select the product that is felt to be most suitable.

SFF meets during the T10 (see www.t10.org) and T11 (see www.t11.org) weeks, and SSWGs (Specific Subject Working Groups) are held at the convenience of the participants. Material presented to SFF becomes public domain, and there are no restrictions on the open mailing of the presented material by Members.

Many of the specifications developed by SFF have either been incorporated into standards or adopted as standards by ANSI, EIA, JEDEC and SAE.

For those who wish to participate in the activities of the SFF TWG, the signup for membership can be found at:

http://www.snia.org/sff/join

The complete list of SFF Specifications which have been completed or are currently being worked on by the SFF Committee can be found at:

http://www.snia.org/sff/specifications/SFF-8000.xls

If you wish to know more about the SFF TWG, the principles which guide the activities can be found at:

http://www.snia.org/sff/specifications/SFF-8032.PDF

Suggestions for improvement of this specification will be welcome, they should be submitted to:

http://www.snia.org/feedback

Change History

Revision 0.5

- Changed SFF-8000.TXT to SFF-8000.xls on page 2 and in section 2.2
- Added Change History
- Added Table of Contents
- Added ASME Y14.5-2009 to section 2.1
- Changed ANSI-Y14.5M to ASME Y14.5-2009 in section 2.3
- Changed the title of Section 4.3 to Pulse Width Jitter

Table of Contents

1. Scope 5
1.1 Copyright 5
1.2 Disclaimer 5

2. References 5
2.1 Industry Documents 5
2.2 Sources 5
2.3 Conventions 6
2.4 Definitions 6

3. Overview 6

4. Physical Layer 7
4.1 Electrical Characteristics 7
4.2 Encoding Scheme 8
4.3 Pulse Width Jitter 8

5. Link Layer 8
5.1 Packet Layout 8
5.2 Checksum Definition 9
5.3 Packet Transmission Behavior 9
5.4 Data Type Definition 10
5.5 Example Packet 10

6. Programming Interface 11

1. Scope

This specification defines a protocol for transmitting digital information by pulse width modulation of the storage device ready/activity signal.

1.1 Copyright

1.2 Disclaimer

Suggestions for revisions should be directed to http://www.snia.org/feedback/

2. References

2.1 Industry Documents

- ASME Y14.5-2009 Dimensioning and Tolerancing

2.2 Sources

There are several projects active within the SFF TWG. The complete list of specifications which have been completed or are still being worked on are listed in http://www.snia.org/sff/specifications/SFF-8000.xls

Copies of ANSI standards may be purchased from the InterNational Committee for Information Technology Standards (http://www.techstreet.com/incitsgate.tmpl).

2.3 Conventions

The dimensioning conventions are described in ASME Y14.5-2009, Geometric Dimensioning and Tolerancing. All dimensions are in millimeters, which are the controlling dimensional units (if inches are supplied, they are for guidance only).

The ISO convention of numbering is used i.e., the thousands and higher multiples are separated by a space and a period is used as the decimal point. This is equivalent to the English/American convention of a comma and a period.

American	French	ISO
0.6	0,6	0.6
1,000	1 000	1 000
1,323,462.9	1 323 462,9	1 323 462.9

2.4 Definitions

For the purpose of SFF Specifications, the following definitions apply:

BMC	Baseboard Management Controller
HDD	Hard Disk Drive
LED	Light Emitting Diode
SAS	Serial Attached SCSI
SATA	Serial ATA
SNIA	Storage Networking Industry Association
SSD	Solid State Drive
TWG	Technical Working Group

3. Overview

In server and storage platforms that support multiple storage device bays (either HDD or SSD) the entity responsible for the management of the chassis may not be the same entity that controls the data path of the storage devices. In many cases chassis management is handled by a Baseboard Management Controller (BMC) that is independent of the main compute complex (as in the case of a monolithic server) or multiple compute nodes (as in the case of a modular chassis). In these cases the BMC requires an out-of-band method to determine the state of each storage device. This method is primarily intended to communicate device temperature information but could be used for other device status information as well.

SFF-8609 is a method to receive device status information by pulse width modulation of the drive ready/activity signal. During periods of activity the drive ready/activity signal oscillates at a rate of 10 Hertz with a 50% duty cycle. This produces a square wave signal with a total period of 100mS (50mS positive pulse and 50mS negative pulse). The idea is to vary the positive and negative pulse widths between 42mS and 58mS so as to embed a nibble of digital data in each pulse. The variation of pulses between 42mS and 58mS is imperceptible to the human eye.

4. Physical Layer

The following sections specify the physical layer characteristics of the protocol.

4.1 Electrical Characteristics

The storage device READY LED signal is described in both the SAS SPL-5 specification and the SATA specification as a signal used to activate an externally visible LED that indicates the state of readiness and activity of the target device. The signal is based on an open-collector or open-drain active low driver and is not suitable for directly driving the LED itself. The LED and current limiting circuitry shall be external to the target device. The figure below show a logical example of an LED driver circuit.

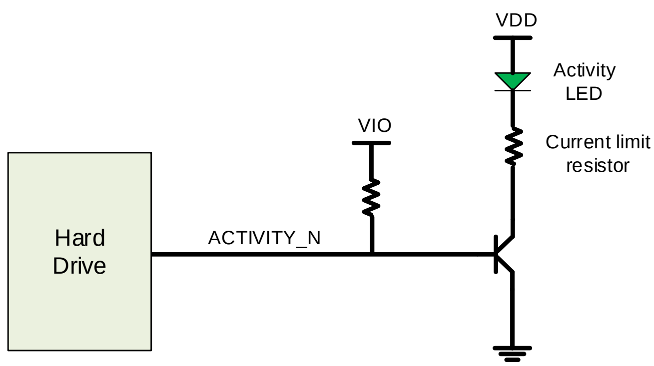

Figure 1

To implement the SFF-8609 protocol additional logic would be added to monitor the ready/activity signal. The additional logic would measure the pulse width of each pulse to determine if digital information is embedded in the signal. The figure below shows a logical example of an implementation:

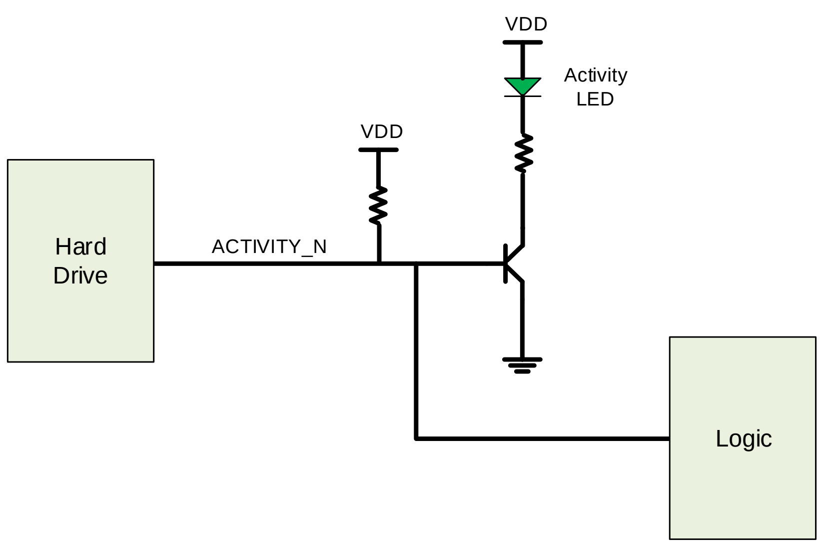

Figure 2

4.2 Encoding Scheme

The storage device shall encode a nibble (4 bits) of digital information in each positive or negative pulse of the ready/activity signal by modulating the width of the pulses. The resolution of the pulse widths shall be 1mS starting at 42mS for a value of 0 hex and extending to 58mS for a value of F hex. A pulse width of 50mS shall represent a null pulse (the nominal pulse width when not transmitting data). The figure below is a graphical representation of the encoding scheme:

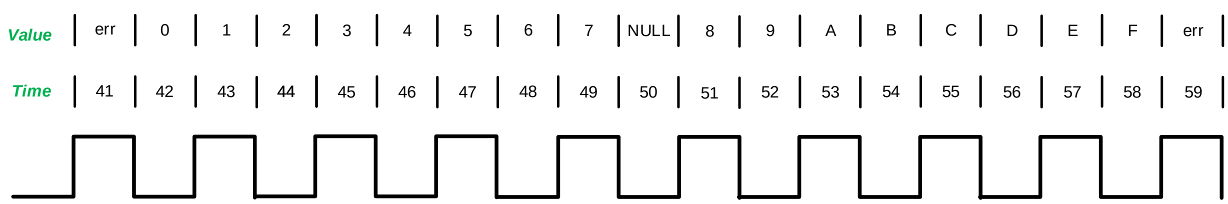

Figure 3

(Note: the time values are in milliseconds)

Pulse widths that are either less than 42mS or greater than 58mS shall be considered errors and the receiver shall ignore the current packet in transmission.

4.3 Pulse Width Jitter

Valid pulse widths shall range between 42mS and 58mS +/- 0.25mS on 1mS boundaries. Pulse widths that are outside of the allowable jitter margin shall be considered errors. In the event of a pulse width error the receiver shall ignore the current packet in transmission.

5. Link Layer

The following sections specify the link layer characteristics of the protocol.

5.1 Packet Layout

A valid packet consists of a total of nine pulses. Each packet is preceded by a null (50mS) pulse and consist of a data type pulse, four data pulses, and three checksum pulses. The figure below shows a graphical representation of a valid packet

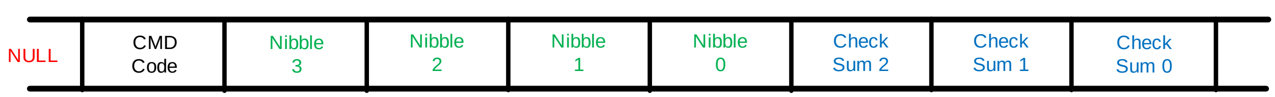

Figure 4

The data code is a single nibble that specifies one of sixteen different data types that the four data nibbles can represent.

The four data nibbles represent the actual data being transmitted. The data nibbles are transmitted most significant nibble first.

The three checksum nibbles are used to calculate a checksum value that is used to validate the data code and data nibbles. The checksum nibbles are transmitted most significant nibble first.

5.2 Checksum Definition

The checksum is calculated using a modified Fletcher checksum algorithm. The modified Fletcher algorithm catches all double pulse errors without significantly increasing the computational complexity or increasing the checksum length.

The following is the algorithm used to calculate the checksum.

```
Initialize variables A and B to 00 hex

A = A + “Data Code”
B = B + A
A = A + “Nibble 3”
B = B + A
A = A + “Nibble 2”
B = B + A
A = A + “Nibble 1”
B = B + A
A = A + “Nibble 0”
B = B + A

Checksum = (B<<4) + A
```

Where << 4 is a binary shift left by 4 operation. This yields a 12 bit checksum value.

5.3 Packet Transmission Behavior

The following specifies the serial transmission behavior for the packets.

- Packets shall always be preceded by a NULL (50 mS) pulse
- Packets shall be in sequence. Receipt of a NULL pulse or an ERROR pulse anywhere within the packet shall cause the receiver to ignore the packet.
- An invalid checksum comparison shall cause the receiver to ignore the packet.

5.4 Data Type Definition

The following table shows the current Data Code definitions:

DATA CODE	DATA LENGTH	CHECKSUM LENGTH	ITEM	DESCRIPTION
NULL	0	0	Null packet	Preamble for receiver sync
x0	4	3	Instantaneous Temperature	Degrees C (note 1)
x1	tbd	tbd	reserved	Reserved for future use
***	***	***	***	***
xD	tbd	tbd	reserved	Reserved for future use
xE	4	3	Protocol Revision Code	Revision control (note 2)
xF	4	3	Stopping transmission	Data nibbles are all x0 (note 3)

Table 1

Data codes x1 through xD are reserved for future use.

Note 1: The data value for instantaneous temperature is an 8 bit two's complement integer value in nibbles 0 and 1. This gives a temperature range of +128° C to -127° C. Nibbles 2 and 3 are reserved for future use and shall be set to 0.

Note 2: A sequence of five Protocol Revision Code packets shall be sent at one second intervals as specified in the protocol standards referencing this SFF document. The value of the data nibbles shall reflect the major.minor revision of the SFF-8609 specification where the major revision is contained in data nibbles 3 & 2 and the minor revision is contained in data nibbles 1 & 0.

Note 3: The use of the stopping transmission packet is specified in the protocol standards referencing this SFF document.

5.5 Example Packet

The following section shows an example of a packet that contains an instantaneous temperature reading of +25° C.

The temporal format of the packet is as follows:

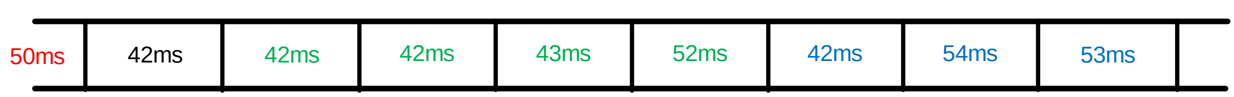

Figure 5

In this case the Data Code is 0x0h and a temperature of +25° C translates to a hexadecimal integer value of 0x0019h. The values in the packet are as follows:

NULL	0	0	0	1	9	0	B	A

Figure 6

The checksum calculation for a Data Code of 0x0h and a data value of 0x0019h is as follows:

Initialize variables A and B to 00 hex

A = A + “Data Code”	00h + 00h = 00h
B = B + A	00h + 00h = 00h
A = A + “Nibble 3”	00h + 00h = 00h
B = B + A	00h + 00h = 00h
A = A + “Nibble 2”	00h + 00h = 00h
B = B + A	00h + 00h = 00h
A = A + “Nibble 1”	00h + 01h = 01h
B = B + A	00h + 01h = 01h
A = A + “Nibble 0”	01h + 09h = 0Ah
B = B + A	01h + 0Ah = 0Bh

Checksum = (B<<4) + A   0B0h + 0Ah = 0BAh

The packet values shown above translate into a pulse width modulated digital bit stream with the following temporal characteristics:

50ms	42ms	42ms	42ms	43ms	52ms	42ms	54ms	53ms	50ms

Figure 7

6. Programming Interface

Definitions for the control of this protocol shall be left to the respective standards groups (SAS and SATA).