# Summary of Modifications

## Instructions for the Editor

1. Clause 4.2.4.1
    - Table 6
        i. Add the following descriptors:
            1. 0001 000Fh Congestion Detection Capability descriptor
            2. 0002 0001h Link Integrity Notification descriptor
            3. 0002 0002h Delivery Notification descriptor
            4. 0002 0003h Peer Congestion Notification descriptor
            5. 0002 0004h Congestion Notification descriptor
            6. 0003 0001h FPIN Registration descriptor
2. Clause 4.3.1
    - Table 9
        i. Add the following ELS codes
            1. 16h Fabric Performance Impact Notification (FPIN)
            2. 19h Register Diagnostic Functions (RDF)
3. Clause 4.3.52.3
    - Remove requirement for Link Fault descriptor under table 180
4. Clause 4.3.52.4
    - Remove requirement for Link Fault descriptor under table 181
5. Clause 4.3.52.5.3
    - Add new clause describing the Congestion Detection Capability descriptor
6. Clause 4.3.53
    - Add new clause describing Register Diagnostic Functions (RDF)
7. Clause 4.3.54
    - Add new clause describing Fabric Performance Impact Notification (FPIN)

*Changes are shown highlighted in yellow.

### 4.2.4 Link Service TLV descriptors

#### 4.2.4.1 Overview

Link Service TLV descriptors are specified in table 6.

**Table 6 – Link Service TLV descriptors**

| Tag value | Descriptor | Reference |
|---|---|---|
| 0000 0000h | Reserved | |
| 0000 0001h | Link Service Request Information | 4.2.4.2 |
| 0000 0002h | Reserved | |
| 0000 0003h | N_Port_ID descriptor | 4.2.4.3 |
| 0001 0000h | SFP Diagnostics descriptor | 4.3.49.5.4 |
| 0001 0001h | Port Speed descriptor | 4.3.49.5.1 |
| 0001 0002h | Link Error Status Block descriptor | 4.3.49.5.2 |
| 0001 0003h | Port Names descriptor | 4.3.49.5.3 |
| 0001 0004h | QSFP Diagnostics descriptor | 4.3.49.5.5 |
| 0001 0005h | FEC Status descriptor | 4.3.49.5.6 |
| 0001 0006h | Buffer Credit descriptor | 4.3.49.5.7 |
| 0001 0007h | Optical Element Data descriptor | 4.3.49.5.8 |
| 0001 0008h | Optical Product Data descriptor | 4.3.49.5.9 |
| 0001 0009h | Priority Range descriptor | 4.3.50.4 |
| 0001 000Ah | VEM ID descriptor | 4.3.51.3 |
| 0001 000Bh | Instantiated VE Mapping descriptor | 4.3.51.3 |
| 0001 000Ch | Deinstantiated VE Mapping descriptor | 4.3.51.3 |
| 0001 000Dh | Link Fault Capability descriptor | 4.3.52.5.2 |
| 0001 000Eh | Reserved | |
| 0001 000Fh | Congestion Detection Capability descriptor | 4.3.52.5.3 |
| 0002 0000h | Reserved | |
| 0002 0001h | Link Integrity Notification descriptor | 4.3.54.7.2 |
| 0002 0002h | Delivery Notification descriptor | 4.3.54.7.3 |
| 0002 0003h | Peer Congestion Notification descriptor | 4.3.54.7.4 |
| 0002 0004h | Congestion Notification descriptor | 4.3.54.7.5 |
| 0003 0000h | Reserved | |

| | | |
|---|---|---|
| 0003 0001h | FPIN Registration descriptor | 4.3.53.5.2 |
| All other values | Reserved | |

### 4.3 Extended Link Service requests

#### 4.3.1 Introduction

A Sequence Initiator shall transmit an ELS Sequence in order to solicit the destination Nx_Port to perform a link-level function or service. Unless otherwise noted, Extended Link Service requests shall not be issued prior to completion of N_Port Login. Table 9 applies to ELSs sent to or received by all valid addresses, including well known addresses. FLOGI is required before any other ELS if a Fabric is present.

The LFA is used as the destination ID (D_ID) in the LINIT and LSTS ELS Request Sequences, and is used as the source ID (S_ID) in the Reply Sequences. No other Sequences shall be directed to a LFA.

**Table 9 – ELS_Command codes**

| Value (Bits 31-24) | Description | Abbr. | Reference | N_Port Login Required | TLV Format |
|---|---|---|---|---|---|
| **01h** | Link Service Reject | LS_RJT | 4.4.4 | N/A | No |
| **02h** | Link Service Accept | LS_ACC | 4.4.2 | N/A | No |
| **03h** | N_Port Login | PLOGI | 4.3.7 | No | No |
| **04h** | F_Port Login | FLOGI | 4.3.7 | No | No |
| **05h** | Logout | LOGO | 4.3.8 | No | No |
| **06h** | Abort Exchange – obsolete | ABTX | N/A | N/A | No |
| **07h** | Read Connection Status – obsolete | RCS | N/A | N/A | No |
| **08h** | Read Exchange Status Block – obsolete | RES | N/A | N/A | No |
| **09h** | Read Sequence Status Block – obsolete | RSS | N/A | N/A | No |
| **0Ah** | Request Sequence Initiative | RSI | 4.3.12 | Yes | No |
| **0Bh** | Establish Streaming | ESTS | 4.3.6 | Yes | No |
| **0Ch** | Estimate Credit | ESTC | 4.3.5 | Yes | No |
| **0Dh** | Advise Credit | ADVC | 4.3.3 | Yes | No |
| **0Eh** | Read Timeout Value | RTV | 4.3.10 | Yes | No |
| **0Fh** | Read Link Error Status Block | RLS | 4.3.9 | Yes | No |
| **10h** | Echo | ECHO | 4.3.4 | No | No |
| **11h** | Test | TEST | 4.3.13 | Yes | No |
| **12h** | Reinstate Recovery Qualifier | RRQ | 4.3.11 | Yes | No |
| **13h** | Read Exchange Concise | REC | 4.3.36 | Yes | No |
| **14h** | Reserved for legacy implementations[a] | | | | |
| **16h** | Fabric Performance Impact Notification | FPIN | 4.3.54 | No | Yes |
| **17h** | Exchange Diagnostic Capabilities[b] | EDC | 4.3.51 | No | Yes |
| **18h** | Read Diagnostic Parameters | RDP | 4.3.48 | Yes | Yes |
| **19h** | Register Diagnostic Functions | RDF | 4.3.53 | No | Yes |
| **20h** | Process Login | PRLI | 4.3.19 | Yes | No |
| **21h** | Process Logout | PRLO | 4.3.20 | Yes | No |
| **23h** | Test Process Login | TPLS | 4.3.21 | Yes | No |

[a] Some early implementations of FCP-2 may have used the value 14h for SRR (Sequence Retransmission Request). This code is permanently reserved in this standard to avoid conflicts with such implementations. See FCP-3 for the standard implementation of SRR as an FC-4 Link Service.

[b] N_Port Login required if the D_ID is the N_Port_ID of an N_Port. N_Port Login not required if the D_ID is the F_Port Controller.

### 4.3.52 Exchange Diagnostic Capabilities (EDC)

#### 4.3.52.1 Description

The EDC ELS is used to exchange diagnostic capabilities parameters between requesting and responding FC_ports.

#### 4.3.52.2 Protocol

a) Exchange Diagnostic Capabilities Request Sequence; and
b) LS_ACC or LS_RJT Reply Sequence

#### 4.3.52.3 Request Sequence

**Addressing:** The S_ID field designates the source Nx_Port requesting the diagnostic parameters. The D_ID field shall be set as follows:

a) to the F_Port Controller (FFFFEh); or
b) to any Nx_Port N_Port_ID

**Payload**: The format of the EDC Request Payload is shown in table 180.

**Table 180 – EDC Payload**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| **0** | EDC (17h) | 00h | 00h | 00h |
| **1** | Descriptor list length ((n-1)*4) bytes) | | | |
| **2 - n** | Diagnostic capability descriptors | | | |

~~The diagnostic capability descriptors (see 4.3.52.5) in an EDC request shall include the Link Fault Capability descriptor.~~

An EDC request shall include zero or more Diagnostic Capability descriptors (see 4.3.52.5).

#### 4.3.52.4 Reply Sequence

**LS_RJT**: LS_RJT signifies rejection of the EDC command.

**LS_ACC**: LS_ACC signifies acceptance of the request and provides the requested data. The format of the LS_ACC Payload is shown in Table 181.

**Table 181 – EDC LS_ACC Payload**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| **0** | LS_ACC (02h) | 00h | 00h | 00h |
| **1** | Descriptor list length ((n-1)*4) bytes) | | | |
| **2** | MSB | Link Service Request Information descriptor (see 4.2.4.2) | | |
| **3** | | | | |
| **4** | | | | LSB |
| **5 - n** | Diagnostic capability descriptors | | | |

~~The diagnostic capability descriptors (see 4.3.52.5) in an EDC LS_ACC shall include the Link Fault Capability descriptor.~~

The diagnostic capability descriptors in an EDC LS_ACC shall include zero or more Diagnostic Capability descriptors (see 4.3.52.5).

**4.3.52.5 Diagnostic capability descriptors**

**4.3.52.5.1 Overview**

The diagnostic capability descriptors included in an EDC request correspond to functions supported by the requesting FC_Port.

The diagnostic capability descriptors included in an EDC response correspond to functions supported by the responding FC_Port.

**4.3.52.5.2 Link Fault Capability descriptor**

Link Degrade Signaling is supported by monitoring errors in the 64GFC Reed-Solomon FEC logic (see FC-FS-5). The Link Fault Capability descriptor provides a mechanism to exchange the Link Degrade Signaling function parameters between two ports. This allows each port to determine the error rate associated with the Link Degrade Signal when it is received.

The Link Fault Capability descriptor is shown in table 182.

**Table 182 – Link Fault Capability descriptor**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| **0** | Link Fault Capability descriptor tag = 0001 000Dh | | | |
| **1** | Descriptor Length (12 bytes) | | | |
| **2** | Degrade Activate Threshold | | | |
| **3** | Degrade Deactivate Threshold | | | |
| **4** | FEC Degrade Interval | | | |

**Degrade Activate Threshold**: The value of the Degrade_Activate_Threshold register (see FC-FS-5).

**Degrade Deactivate Threshold**: The value of the Degrade_Deactivate_Threshold register (see FC-FS-5).

**FEC Degrade Interval**: The value of the FEC_Degrade_interval register (see FC-FS-5).

**4.3.52.5.3 Congestion Detection Capability descriptor**

Congestion Detection Signaling is supported by sending and/or receiving the Warning Congestion Signal and Alarm Congestion Signal (see FC-FS-6). The Congestion Detection Capability descriptor provides a mechanism to exchange the Congestion Detection Signaling function parameters between two ports. The format of the Congestion Detection Capabilities descriptor is shown in Table 183.

**Table 183 - Congestion Detection Capabilities descriptor**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | Congestion Detection Capability Descriptor tag = 0001 000Fh | | | |
| 1 | Descriptor Length (16 bytes) | | | |
| 2 | Transmit Signal Capability | | | |
| 3 | Transmit Signal Frequency | | | |
| 4 | Receive Signal Capability | | | |
| 5 | Receive Signal Frequency | | | |

**Transmit Signal Capability**: Defines the ability of the FC_Port to transmit the primitives used for Congestion Signals (see FC-FS-6). An FC_Port that is capable of transmitting only one signal value shall indicate transmission of the Warning Congestion Signal. The definition of the Transmit Signal Capability methods is listed in Table 184.

**Table 184 – Signal Capability description**

| Bit | Description |
|---|---|
| 31 – 4 | Reserved |
| 3 – 0 | **Signal Capability values**:<br>0 = The Warning Congestion Signal and the Alarm Congestion Signal are not supported<br>1 = The Warning Congestion Signal is supported<br>2 = The Warning Congestion Signal and the Alarm Congestion Signal are supported<br>All other values are reserved |

**Transmit Signal Frequency**: The Transmit Signal Frequency describes the minimum time between transmissions of Congestion Signals. The Transmit Signal Frequency fields are defined in Table 185.

**Table 185 – Signal Frequency description**

| Bit | Description |
|---|---|
| 31 – 26 | Reserved |
| 25 – 16 | **Signal Frequency Count**:<br>0 = Reserved<br>1 to 999 = The time between signals in the units indicated in the Signal Frequency Units field.<br>All other values are reserved |
| 15 - 4 | Reserved |
| 3 - 0 | **Signal Frequency Units**:<br>0 = Reserved<br>1 = Seconds<br>2 = Milliseconds<br>All other values are reserved |

**Receive Signal Capability**: Defines the ability of the FC_Port to receive the primitives used for Congestion Signals (see FC-FS-6). An FC_Port that is capable of receiving only one signal value shall indicate reception of the Warning Congestion Signal. The definition of the Receive Signal Capability methods is listed in Table 184.

**Receive Signal Frequency**: The Transmit Signal Frequency describes the minimum time between transmissions of Congestion Signals. The Receive Signal Frequency describes the minimum time between receiving of Congestion Signals. The Receive Signal Frequency fields are defined in Table 185.

At the completion of the EDC exchange, each end-point of the link adopts the least capable settings for Signal Capability and Signal Frequency of the two end-points thus, a port's frequency of transmission is the larger of that port's transmit value and the connected port's receive value.

> Example #1, the EDC sender is capable of transmitting and receiving the Warning Congestion Signal and the Alarm Congestion Signal at a rate of 10-milliseconds. The EDC responder is capable of transmitting and receiving only the Warning Congestion Signal at a rate of 100-milliseconds. At the conclusion of the EDC exchange, both end-points adopt transmission and reception of the Warning Congestion Signal at a rate of 100-milliseconds.
>
> Example #2, the EDC sender is capable of transmitting and receiving the Warning Congestion Signal and the Alarm Congestion Signal at a rate of 10-milliseconds. The EDC responder is capable of transmitting the Warning Congestion Signal at a rate of 100-milliseconds and is capable of receiving the Warning Congestion Signal and the Alarm Congestion Signal at a rate of 10-milliseconds. At the conclusion of the EDC exchange, the EDC sender adopts transmitting the Warning Congestion Signal and the Alarm Congestion Signal at a rate of 10-milliseconds and adopts receiving the Warning Congestion Signal at a rate of 100-milliseconds; likewise, the EDC responder adopts receiving the Warning Congestion Signal and the Alarm Congestion Signal at a rate of 10-milliseconds and adopts transmitting the Warning Congestion Signal at a rate of 100-milliseconds.

### 4.3.53 Register Diagnostic Functions (RDF)

#### 4.3.53.1 Description

The RDF ELS requests the Fabric Controller to add the Nx_Port that is sending the RDF Request to the list of Nx_Ports registered to receive the specified FPIN ELS descriptors.

#### 4.3.53.2 Protocol

a) Register Diagnostic Function Request Sequence; and
b) LS_ACC or LS_RJT Reply Sequence

#### 4.3.53.3 Request Sequence

**Addressing**: The S_ID field designates the source Nx_Port requesting registration of the diagnostic functions. The D_ID designates the Fabric Controller, FFFFFDh.

**Payload**: The format of the RDF Request Payload is shown in table 186.

**Table 186 – RDF ELS Request format**

| Bits / Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | RDF (19h) | 00h | 00h | 00h |
| 1 | Descriptor list length ((n-1)*4) bytes) | | | |
| 2 - n | Descriptor list | | | |

**Descriptor list:** The list of descriptors identifying the diagnostic functions supported by requesting Nx_Port (see 4.3.53.5).

#### 4.3.53.4 Reply Sequence

**LS_RJT**: LS_RJT signifies rejection of the RDF command.

**LS_ACC**: LS_ACC signifies acceptance of the RDF Request and registration of the Nx_Port to receive the specified FPIN ELS descriptors. The format of the LS_ACC Payload is shown in Table 187.

**Table 187 – RDF ELS Reply format**

| Bits / Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | LS_ACC (02h) | 00h | 00h | 00h |
| 1 | Descriptor list length ((n-1)*4) bytes) | | | |
| 2 | MSB | Link Service Request Information descriptor (see 4.2.4.2) | | |
| 3 | | | | |
| 4 | | | | LSB |
| 5 - n | Descriptor list | | | |

**Descriptor list:** The list of descriptors describing the diagnostic functions supported by the requesting Nx_Port (see 4.3.53.5).

The RDF command provides a method for Nx_Ports to indicate to the Fabric Controller the specific diagnostics functions supported by the FC_Port and the range of diagnostic capabilities associated with each of the supported functions.

Prior to registration, the FPIN ELS shall not be transmitted to the Nx_Port by the Fabric Controller. After registration, the FPIN ELS transmitted to the Nx_Port shall contain only the diagnostic functions successfully registered in the RDF ELS request/response exchange.

#### 4.3.53.5 Diagnostic function descriptors

##### 4.3.53.5.1 Overview

The diagnostic function descriptors included in an RDF Request correspond to functions supported by the requesting Nx_Port and to be registered with the Fabric Controller.

The diagnostic function descriptors included in an RDF LS_ACC correspond to functions supported by the Fabric Controller and the requesting Nx_Port. The RDF LS_ACC from the Fabric Controller indicates the Nx_Port is registered to receive the FPIN ELS.

##### 4.3.53.5.2 FPIN Registration descriptor

The FPIN Registration descriptor registers an Nx_Port with the Fabric Controller to receive FPIN ELS. The format of the FPIN Registration descriptor is shown in Table 188.

**Table 188 – FPIN Registration descriptor**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | FPIN Registration descriptor tag = 0003 0001h | | | |
| 1 | FPIN Registration descriptor length ((n-1)*4) bytes) | | | |
| 2 | Count | | | |
| 3-n | Descriptor Tag list | | | |

**Count**: Number of descriptor tags indicating the FPIN functions supported.

**Descriptor Tag list:** List of FPIN descriptor tags supported. The Descriptor Tag list in the RDF Request indicates the functions supported by the Nx_Port. The Descriptor Tag list in the RDF LS_ACC indicates the functions supported by the Nx_Port and the Fabric Controller, which may be a subset of the functions listed in the RDF Request.

### 4.3.54 Fabric Performance Impact Notifications (FPIN)

#### 4.3.52.1 Description

The FPIN ELS notifies registered Nx_Ports of Fabric events or behaviors in the Fabric. The FPIN ELS contains a list of notification descriptors for each detected event, which may be coalesced for simultaneously detected events. Each descriptor contains a description of the Fabric event.

#### 4.3.54.2 FPINs issued by the Fabric Controller

The Fabric Controller may issue an FPIN Request to registered Nx_Ports when the Fabric detects an event. The distribution scope of FPIN Requests is defined by the Fabric notification descriptor (see 4.3.54.7) and is limited to the Nx_Ports in the zone membership list of the Nx_Port associated with the event. An event may include any of the following:

a) Link Integrity event;
b) Delivery event;
c) Peer Congestion event; or
d) Congestion event.

#### 4.3.54.3 FPINs issued by the Nx_Port

An Nx_Port may issue an FPIN ELS Request to the Fabric Controller when an event is detected by an Nx_Port. Fabric Controllers are implicitly registered to receive FPINs after a successful implicit or explicit FLOGI. FPIN ELS Requests received by the Fabric Controller from an Nx_Port are forwarded by the Fabric Controller. The forwarding scope of the FPIN Request is defined by the Fabric notification descriptor (see 4.3.54.7) and is limited to the Nx_Ports in the zone membership list of the requesting Nx_Port. An event may include any of the following:

a) Link Integrity event; or
b) Congestion event.

#### 4.3.54.4 Protocol

a) Fabric Performance Impact Notification Request Sequence
b) None or LS_RJT Reply Sequence

#### 4.3.54.5 Request Sequence

**Addressing:** If the Fabric is using FPIN to notify a registered Nx_Port of a Fabric event (see 4.3.54.2), the S_ID is the Fabric Controller (FFFFFDh) and the D_ID is the address of the registered Nx_Port destination. If an Nx_Port is using FPIN to notify the Fabric of a Fabric event (see 4.3.54.3), the S_ID designates the Nx_Port indicating a Fabric event and the D_ID is the Fabric Controller (FFFFFDh).

**Payload**: The format of the FPIN Request Payload is shown in table 189.

**Table 189 – FPIN ELS format**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | FPIN (16h) | 00h | 00h | 00h |
| 1 | Descriptor list length ((n-1)*4) bytes) | | | |
| 2 - n | Descriptor list | | | |

**Descriptor list:** The list of notification descriptors identifying the event detected by the Fabric Controller or an Nx_Port (see 4.3.54.7).

**4.3.54.6 Reply Sequence**

**LS_RJT**: LS_RJT signifies rejection of the FPIN command.

**LS_ACC**: LS_ACC is not required.

**4.3.54.7 Fabric notification descriptors**

**4.3.54.7.1 Overview**

The Fabric notification descriptors included in an FPIN request from the Fabric Controller correspond to events detected by the Fabric (see 4.3.54.2).

The Fabric notification descriptors included in an FPIN request from an Nx_Port correspond to events detected by the Nx_Port (see 4.3.54.3).

**4.3.54.7.3 Link Integrity notification descriptor**

The Link Integrity notification descriptor indicates an error threshold has been exceeded. The format of the Link Integrity notification descriptor is shown in Table 190.

**Table 190 – Link Integrity notification descriptor**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | Link Integrity notification descriptor tag = 0002 0001h | | | |
| 1 | Descriptor length ((n-1)*4) bytes) | | | |
| 2 | MSB — Detecting Port Name | | | |
| 3 | | | | LSB |
| 4 | MSB — Attached Port Name | | | |
| 5 | | | | LSB |
| 6 | Event Type | | Event Modifier | |
| 7 | Event Threshold | | | |
| 8 | Event Count | | | |
| 9 | Port Name Count | | | |
| 10 - n | Port Name List | | | |

**Detecting Port Name**: The Name_Identifier of the FC_Port detecting the notification event.

**Attached Port Name**: The Name_Identifier of the FC_Port connected to the FC_Port detecting the notification event.

**Event Type**: The type of event reported in the Link Integrity notification descriptor. The Event Type values are show in Table 191.

**Table 191 – Link Integrity notification Event Type values**

| Value (hex) | Description |
|---|---|
| **00h** | Unknown |
| **01h** | Link Failure |
| **02h** | Loss-of-Synchronization |
| **03h** | Loss-of-Signal |
| **04h** | Primitive Sequence Protocol Error |
| **05h** | Invalid Transmission Word |
| **06h** | Invalid CRC |
| **0Fh** | Device Specific |
| All other values are reserved | |

**Event Modifier:** An implementation specific value describing the Event Type (I.e., information describing the Device Specific Event Type).

**Event Threshold**: The duration in milliseconds of the Link Integrity detection cycle.

**Event Count**: The minimum number of event occurrences during the Event Threshold to cause generation of a Link Integrity event.

**Port Name Count**: Number of Port Names in the Port Name List. If the attached FC_Port is the only accessible port then the Count is set to zero and the Port Name List is empty.

**Port Name List**: The list of N_Port_Names accessible through the attached ported. An N_Port_Name is the Name_Identifier associated with the port (see FC-FS-6).

Link Integrity events detected by an F_Port are distributed by the Fabric Controller to all the Nx_Ports in the zone membership list(s) associated with the PN_Port attached to the detecting F_Port.

Link Integrity events detected by an Nx_Port are sent to the Fabric Controller and distributed by the Fabric Controller to all the Nx_Ports in the zone membership list(s) associated with the PN_Port of the detecting Nx_Port.

Link Integrity events detected by an E_Port are distributed by the Fabric Controller to all the Nx_Ports affected by the event.

#### 4.3.54.7.4 Delivery notification descriptor

The Delivery notification descriptor indicates a frame has been discarded by the Fabric. The format of the Delivery notification descriptor is shown in Table 192.

**Table 192 – Delivery notification descriptor**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | Delivery notification descriptor tag = 0002 0002h | | | |
| 1 | Descriptor length (44 bytes) | | | |
| 2 | MSB — Detecting Port Name | | | |
| 3 | | | | LSB |
| 4 | MSB — Attached Port Name | | | |
| 5 | | | | LSB |
| 6 | Delivery Reason Code | | | |
| 7 – 12 | Delivery Event Data (discarded frame header) | | | |

**Detecting Port Name**: The Name_Identifier of the FC_Port detecting the notification event.

**Attached Port Name**: The Name_Identifier of the FC_Port connected to the detecting FC_Port.

**Delivery Reason Code**: The reason for the delivery failure. The Reason Code values are show in Table 193:

**Table 193 – Delivery Reason Code values**

| Value (hex) | Description |
|---|---|
| 00h | Unknown |
| 01h | Timeout |
| 02h | Unable to route |
| 0Fh | Device Specific |
| All other values are reserved | |

**Delivery Event Data:** The event data for a Delivery notification event is the 24-byte frame header of the discarded frame (see FC-FS-6).

Delivery notification events are sent by the Fabric Controller to the FC_Port identified in the S_ID of the discarded frame header.

Delivery notification events are sent by the Fabric Controller on a best effort basis and only for command, response, status, or last frame in sequence frames.

#### 4.3.54.7.5 Peer Congestion notification descriptor

The Peer Congestion notification descriptor indicates a congestion condition has been detected in the Fabric. The format of the Peer Congestion notification descriptor is shown in Table 194.

**Table 194 – Peer Congestion notification descriptor**

| Bits / Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | Peer Congestion notification descriptor tag = 0002 0003h | | | |
| 1 | Descriptor length ((n-1)*4) bytes) | | | |
| 2 | MSB — Detecting Port Name — | | | |
| 3 | | | | LSB |
| 4 | MSB — Attached Port Name — | | | |
| 5 | | | | LSB |
| 6 | Event Type | | Event Modifier | |
| 7 | Event Period | | | |
| 8 | Port Name Count | | | |
| 9 - n | Port Name List | | | |

**Detecting Port Name**: The Name_Identifier of the FC_Port detecting the notification event.

**Attached Port Name**: The Name_Identifier of the FC_Port connected to the detecting FC_Port.

**Event Type**: The type of congestion event detected. The Congestion Event Type values are show in Table 195:

**Table 195 – Congestion Event Type values**

| Value (hex) | Description |
|---|---|
| 00h | Clear/None |
| 01h | Lost Credit |
| 02h | Credit Stall |
| 03h | Oversubscription |
| 0Fh | Device Specific |

| All other values are reserved |
|---|

**Event Modifier:** An implementation specific value describing the Event Type (I.e., information describing the Device Specific Event Type).

**Event Period**: The duration in milliseconds of the detected congestion event.

**Port Name Count**: Number of Port Names in the Port Name List. If the attached FC_Port is the only accessible FC_Port then the Count is set to zero and the Port Name List is empty.

**Port Name List**: The list of N_Port_Names accessible through the attached FC_Ported. An N_Port_Name is the Name_Identifier associated with the FC_Port (see FC-FS-6).

Peer Congestion events detected by an F_Port are distributed by the Fabric Controller to the attached FC_Port and to the Nx_Ports in the zone membership list(s) associated with the PN_Port attached to the detecting F_Port.

Peer Congestion events are sent while the congestion condition persists at the interval indicated by the Event Period.

The status of the Peer Congestion event is cleared when:

a) Peer Congestion events cease for a period of two-times the Event Period,
b) a Peer Congestion event with a Congestion Event Type value of “Clear/None” is received, or
c) the Attached Port Name is no longer in the Fabric (I.e., an RSCN is received indicating the Nx_Port with the Attached Port Name is no longer reachable).

#### 4.3.54.7.6 Congestion notification descriptor

The Congestion notification descriptor indicates a congestion condition has been detected at the FC_Port. The format of the Congestion notification descriptor is show in Table 196.

**Table 196 – Congestion notification descriptor**

| Bits<br>Word | 31 .. 24 | 23 .. 16 | 15 .. 08 | 07 .. 00 |
|---|---|---|---|---|
| 0 | Congestion notification descriptor tag = 0002 0004h | | | |
| 1 | Descriptor length (12 bytes) | | | |
| 2 | Event Type | | Event Modifier | |
| 3 | Event Period | | | |
| 4 | Severity | Reserved | | |

**Event Type**: The type of the detected congestion event (see Table 195).

**Event Modifier:** An implementation specific value describing the Event Type (I.e., information describing the Device Specific Event Type).

**Event Period**: The duration in milliseconds of the detected congestion event.

**Severity**: The Status contains a value indicating the distress level of the detected congestion event. The Severity values are show in Table 197:

**Table 197 – Congestion notification Severity values**

| Value (hex) | Description |
|---|---|
| F1h | Warning |
| F7h | Alarm |
| All other values are reserved | |

Congestion events detected by an F_Port are distributed by the Fabric Controller to the attached FC_Port and the Congestion notification descriptor is the only descriptor included in the FPIN ELS payload.

Congestion events are sent to the attached FC_Port while the congestion condition persists at the interval indicated by the Event Period.

The status of the Congestion event is cleared when:

a) Congestion events cease for a period of two-times the Event Period,
b) a Congestion event with a Congestion Event Type value of “Clear/None” is received, or
c) the attached FC_Port is no longer in the Fabric.