

To: INCITS Technical Committee T11
From: David L.Black, EMC Corporation
Email: black_david@emc.com
Date: June 6, 2010
Subject: FC-BB-6: ASFC for the FC pseudowire

Version history:

Revision 0 (6 June 2010) - First version.

Overview:

Changes to the IETF approach to standardizing its portion of the FC pseudo-wire (PW) have removed the Selective Retransmission (SR) protocol that provided retransmission and rate-based flow control. It is necessary to replace SR's flow control functionality to ensure correct FC behavior, but rate-based flow control is not required.

This proposal replaces SR's flow-control functionality by adapting the ASFC (Alternate Simple Flow Control) protocol from FC GFPT to the FC pseudowire.

Removed text is in ~~red strikethrough~~, added text is in blue underline. **Editor's Notes** are in red and indented.

Proposed Changes:

Editor's Note: The current IETF reference for the FC-PW Internet-Draft is draft-ietf-fc-encap-10.txt . Please update the reference to that version and ensure that all fc-encap draft references are to version -10.

4.4.6 Flow control

FC-BB_IP devices shall ensure that TCP flow control and error recovery acts in proper concert with the Fibre Channel BB_Credit flow control mechanism.

Flow control at E_Ports, F_Ports, VE_Ports, and B_Ports shall operate as defined in FC-SW-5.

The Alternate Simple Flow Control (ASFC) mechanism (see 6.4.4) shall be used between FC-BB_GFPT devices and between FC-BB_PW devices. Flow control on FC_Port-facing links is specified in clause 6.

Flow control at FC-BB_GFPT and FC-BB_PW physical interfaces to attached FC devices shall operate as defined in FC-SW-5 or FC-FS-3, as appropriate.

~~Flow control between two FC-BB_PW devices shall operate as specified in draft-ietf-pwe3-fc-encap-09.~~

Editor's Note: Next change is in 6.4.4

6.4.4 Port initialization parameter observation and modification

Transparent FC-BB devices shall note the ISL Flow Control Mode value on inbound (see 3.3.8) and outbound (see 3.3.10) ELP request/reply SW_ACC frames of open LEMs. When the LEM is successfully closed and flow control other than either R_RDY flow control or VC_RDY flow control has been established between the interconnected FC_Ports, the Transparent FC-BB devices behavior is outside the scope of this standard.

If the use of R_RDY or VC_RDY flow control has been negotiated between the interconnected FC_Ports, the FC-BB_GFPT devices shall operate using unchannelized ASFC (see 6.4.7) and FC-BB_PW devices shall operate using [FC-BB_PW ASFC \(see 6.4.7\)](#) ~~Selective Retransmission mode as specified in draft ietf-pwe3-fe-encap-10~~.

Editor's Note: Next Change is in 6.4.7

6.4.7 Transparent FC-BB flow control

6.4.7.1 Overview

Alternate Simple Flow Control (ASFC) (~~see 6.4.7.2~~) shall be supported between FC-BB_GFPT devices (~~see 6.4.7.2~~) and between FC-BB_PW devices (~~see 6.4.7.3~~).

~~FC-BB_PW devices shall support Selective Retransmission as specified in draft ietf-pwe3-fc-encap-40.~~

6.4.7.2 FC-BB_GFPT Alternate Simple Flow Control (ASFC)

ASFC defines two Primitive Signals, ASFC_PAUSE and ASFC_RESUME (see table 16), that may be sent by one FC-BB_GFPT device to a remote FC-BB_GFPT device.

Table 16 specifies the Ordered Set values for the ASFC_PAUSE and ASFC_RESUME FC-BB_GFPT WAN Primitive Signals.

Table 16 – Values of FC-BB_GFPT ASFC_PAUSE and ASFC_RESUME Primitive Signals

FC-BB_GFPT Primitive Signal	Ordered Set (Value)
ASFC_PAUSE	K28.1 D00.0 D00.0 D00.0 - unchannelized
ASFC_RESUME	K28.1 D04.0 D00.0 D00.0 - unchannelized

An FC-BB_GFPT device sends ASFC_PAUSE to a remote FC-BB_GFPT device to direct that device to cease forwarding FC frames and non-GFPT_WAN Primitive Signals. Reception of a single ASFC_PAUSE Primitive Signal suffices to trigger such cessation, which shall take effect:

- a) within 25 microseconds of reception of the ASFC_PAUSE Primitive Signal; or
- b) immediately upon completion of the forwarding of a frame or Primitive Signal if forwarding of a frame or non-GFPT_WAN Primitive Signal is in progress when the ASFC_PAUSE Primitive Signal is received and the time required for such completion is greater than 25 microseconds.

When FC frame and non-GFPT_WAN Primitive Signal forwarding has been suspended due to reception of an ASFC_PAUSE Primitive Signal, it shall remain suspended until at least one ASFC_RESUME Primitive Signal has been received. FC-BB_GFPT devices shall:

- a) not insert ASFC_PAUSE or ASFC_RESUME Primitive Signals within any FC Order Set;
- b) insert ASFC_PAUSE or ASFC_RESUME Primitive Signals on transmission word boundaries; and
- c) not insert more than two ASFC_PAUSE and ASFC_RESUME Primitive Signals within the boundaries of any frame.

Otherwise, ASFC_PAUSE and ASFC_RESUME Primitive Signals are sent at times, in numbers, and at frequencies at the discretion of the sending FC-BB_GFPT device.

Editor's Note: No change to 6.4.7.2 - it's included here for comparison with the new 6.4.7.3.

6.4.7.3 FC-BB_PW Alternate Simple Flow Control (ASFC)

ASFC defines two control frames, ASFC_PAUSE and ASFC_RESUME (see XXX), that may be sent by one FC-BB_PW device to a remote FC-BB_PW device.

Figure X specifies the control frame formats for ASFC_PAUSE and ASFC_RESUME..




Figure X – FC-BB_PW error indication control frame format

Table Y specifies the control frame payloads for ASFC_PAUSE and ASFC_RESUME.

Table Y – FC-BB_PW ASFC_PAUSE and ASFC_RESUME control frame payload values.

<u>Field</u>	<u>Value</u>	<u>Description</u>
<u>Type</u>	<u>0x41</u> <u>0x42</u>	<u>ASFC_PAUSE</u> <u>ASFC_RESUME</u>
<u>PAD</u>	<u>Zeroes</u>	<u>Set to all zeroes.</u>

This standard specifies the use of Type field value 0x11 for transmission of error indications (see 6.4.8.2.2) and specifies the use of Type field values 0x81 and 0x82 for transmission of FC-BB_PW PING and PING_ACK (see 6.4.7.4.3).

An FC-BB_PW device sends ASFC_PAUSE to a remote FC-BB_PW device to direct that device to cease forwarding FC frames and FC Primitive Signals. Reception of a single ASFC_PAUSE control frame suffices to trigger such cessation, which shall take effect:

- a) within 25 microseconds of reception of the ASFC_PAUSE Primitive Signal; or
- b) immediately upon completion of the forwarding of a frame or Primitive Signal if forwarding of a frame or Primitive Signal is in progress when the ASFC_PAUSE Primitive Signal is received and the time required for such completion is greater than 25 microseconds.

Editor's Note: Change 25 microseconds to something larger? FC-BB_PW devices have considerably more functionality and buffering than FC-BB_GFPT devices. 1 millisecond should be more than adequate, and may allow item b) above to be deleted, although keeping it does not cause any harm.

When FC frame and FC Primitive Signal forwarding has been suspended due to reception of an ASFC_PAUSE control frame, it shall remain suspended until at least one ASFC_RESUME control frame has been received. FC-BB_PW devices may send ASFC_PAUSE and ASFC_RESUME control frames at times, in numbers, and at frequencies at the discretion of the sending FC-BB_PW device.

6.4.7.4 PING and PING_ACK signals

Editor’s Note: Renumber 6.4.7.3 and subsections to 6.4.7.4 and subsections

6.4.7.4.1 Overview

To assist in buffer management, as well as for other purposes, it may be useful to measure the latency on the WAN link. To facilitate this, two additional Transparent FC-BB signals are defined, PING and PING_ACK. For FC-BB_GFPT devices these signals are defined as Primitive Signals (see 6.4.7.4.2). For FC-BB_PW devices these signals are defined as control frames (see 6.4.7.4.3). As well as facilitating latency measurements, PING and PING_ACK signals are also used to communicate selected operational information between Transparent FC-BB devices.

The PING signal may be sent at any time by either Transparent FC-BB device, but not more frequently than once every 100 milliseconds. A single PING_ACK signal shall be transmitted by an Transparent FC-BB device in response to each received PING signal, within 25 microseconds of reception of the latter. FC-BB_GFPT devices shall respect FC Ordered Set delineation when inserting PING and PING_ACK Primitive Signals.

Transparent FC-BB devices shall begin transmission of PING signals as soon as possible upon establishment of WAN connectivity in state S1, and shall not exit state S1 until at least one PING_ACK signal has been received from remote Transparent FC-BB devices. This ensures that communication of support for channelized ASFC capabilities takes place between FC-BB_GFPT devices before they exit state S1. Transmission of PING signals is otherwise undertaken at the discretion of individual Transparent FC-BB devices, within the frequency limits specified above.

6.4.7.4.2 FC-BB_GFPT PING and PING_ACK

Table 1 specifies the Ordered Set values for the PING and PING_ACK FC-BB_GFPT WAN Primitive Signals.

Table 1 – FC-BB_GFPT PING and PING_ACK Primitive Signal values

Primitive Signal	Ordered Set (Value)
PING	K28.1 D16.0 D(XNCCCRRRb) D(YZZZZZZZb)
PING_ACK	K28.1 D00.4 D(XNCCCRRRb) D(YZZZZZZZb)

The nomenclature D(MMMMMMMMb), MSB left to LSB right, is interpreted as the 8B/10B D-character equivalent of the binary value MMMMMMMMb.

The X bit indicates support (X=1b) or lack of support (X=0b) of channelized ASFC capability in the FC-BB_GFPT device transmitting the Primitive Signal. Since channelized ASFC is not defined by this standard, the value of the X bit shall be ignored by FC-BB_GFPT devices receiving PING and PING_ACK Primitive Signals.

The N bit may be set to 0b or 1b by devices transmitting PING, and the PING_ACK generated in response shall use the same value of N received in the associated PING.

Table 2 specifies the CCC bit field values.

Table 2 – FC-BB_GFPT PING and PING_ACK CCC bit field values

Value	Description	Reference
000b	No compression	

Table 2 – FC-BB_GFPT PING and PING_ACK CCC bit field values

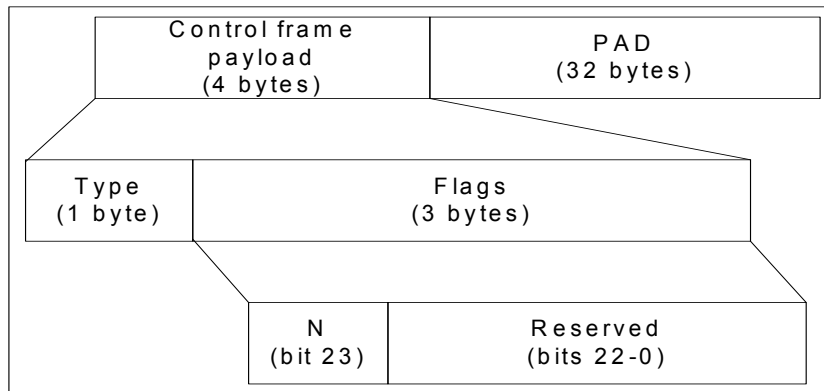
001b	LZS compression	6.4.10.3
010b-111b	Reserved	

The values of the R bits are reserved and shall be set to 0b.

The Y bit specifies vendor-specific use (i.e., Y=1b), or standard use (i.e., Y=0b), of the remaining bits (i.e., bits denoted Z) of the Primitive Signal. If Y=0b, the values of the bits denoted Z are reserved and shall be set to 0b. If Y=1b, the values assigned of the bits denoted Z are vendor specific.

6.4.7.4.3 FC-BB_PW PING and PING_ACK

The FC-BB_PW PING and PING_ACK control frame format is specified in figure 23.

**Figure 23 – FC-BB_PW PING and PING_ACK control frame format**

The FC-BB_PW PING and PING_ACK control frame payload shall be set as specified in table 19.

Table 19 – FC-BB_PW PING and PING_ACK control frame payload values

Field	Value	Description
Type	0x81 0x82	PING PING_ACK
N	0b or 1b	The N bit is set to alternating 0b and 1b. For PING_ACK the N bit is set to the same value as the associated PING.
PAD	Zeroes	Set to all zeroes.

This standard specifies the use of Type field value 0x11 for transmission of error indications (see 6.4.8.2.2) [and specifies the use of Type field values 0x41 and 0x42 for ASFC_PAUSE and ASFC_RESUME \(see 6.4.7.3\).](#)

The FC-BB_PW PING and PING_ACK control frames are encapsulated and formatted for transmission as specified in draft-ietf-pwe3-fc-encap-10.

Editor's Note: Next change is in 6.4.8.2.2

6.4.8.2.2 FC-BB_PW error indication control frame

The FC-BB_PW error indication control frame format is specified in figure 24.




Figure 24 – FC-BB_PW error indication control frame format

The FC-BB_PW error indication control frame payload shall be set as specified in table 20.

Table 20 – FC-BB_PW error indication control frame payload values

Field	Value	Description
Type	0x11	FC-ERROR - indicates a received character error or Loss of Synchronization.
PAD	Zeroes	Set to all zeroes.

This standard specifies the use of Type field values 0x81 and 0x82 for transmission of FC-BB_PW PING and PING_ACK (see 6.4.7.4.3) [and specifies the use of Type field values 0x41 and 0x42 for ASFC_PAUSE and ASFC_RESUME \(see 6.4.7.3\).](#)

The FC-BB_PW error indication control frame is encapsulated and formatted for transmission as specified in draft-ietf-pwe3-fc-encap-10.