Will You Still Love Me When I Turn 64GFC

Live Webcast

Dec. 11, 2018

10:00am PT // 1:00pm ET

About the “When I’m Sixty-Four” lyrics

Will you still need me, will you still feed me
When I'm sixty-four

You'll be older too
And if you say the word
I could stay with you
.
.
Who could ask for more

Remastered Beatles White Album is available for the 50 year anniversary

Today's Speakers

Dean Wallace
Marvell Technology Group
and T11.2 Vice Chairman

deanw@marvell.com

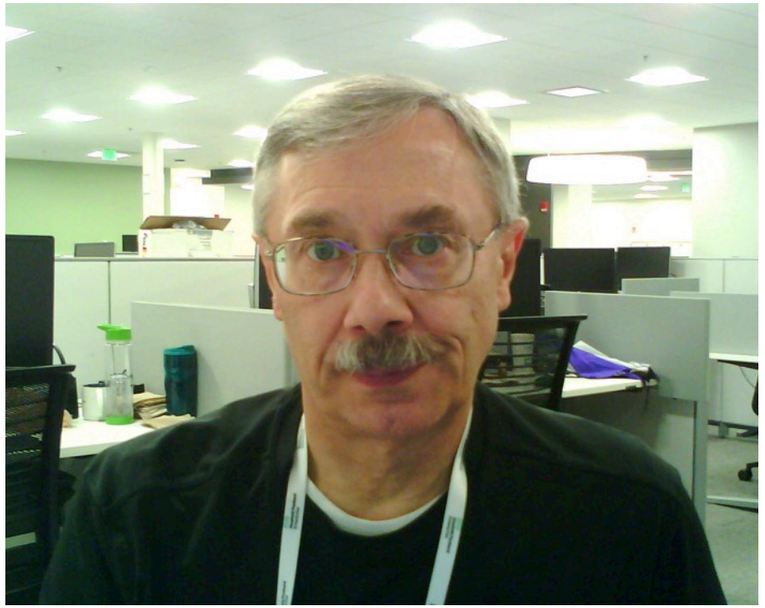

Barry Maskas
HPE Storage Network Consultant
and HPE's Principal T11 member

barry.maskas@hpe.com

Agenda

- Fibre channel terminology and nomenclature overview
- 64GFC standard FC-PI-7
  - Marketing requirements that were used as a drive to goal for T11.2 and T11.3 standards committees
  - Modulation changes for 64GFC
- Optical transceivers used in 64GFC
- Protocol changes required for 64GFC implementation
- 256GFC standard FC-PI-7P
- Future roadmap

Fibre Channel Standards

- A short tour through the acronym soup that are Fibre Channel standards
- The Fibre channel standards focused on in this presentation are:
  - Physical: Fibre-Channel-Physical-Interface, aka FC-PI
  - Protocol: Fibre-Channel-Framing-Signaling, aka FC-FS-5 and FC-FS-4
- A number is appended to the acronym to represent the speed contained in the standard, FC-PI-7 represents 64GFC

Document	Represents…
FC-PI	1GFC 2GFC 4GFC
FC-PI-2	4GFC
FC-PI-4	8GFC
FC-PI-5	16GFC
FC-PI-6 FC-PI-6P	32GFC 128GFC (parallel)
FC-PI-7 **FC-PI-7P**	**64GFC** **256GFC (parallel)**

Fibre Channel Layers

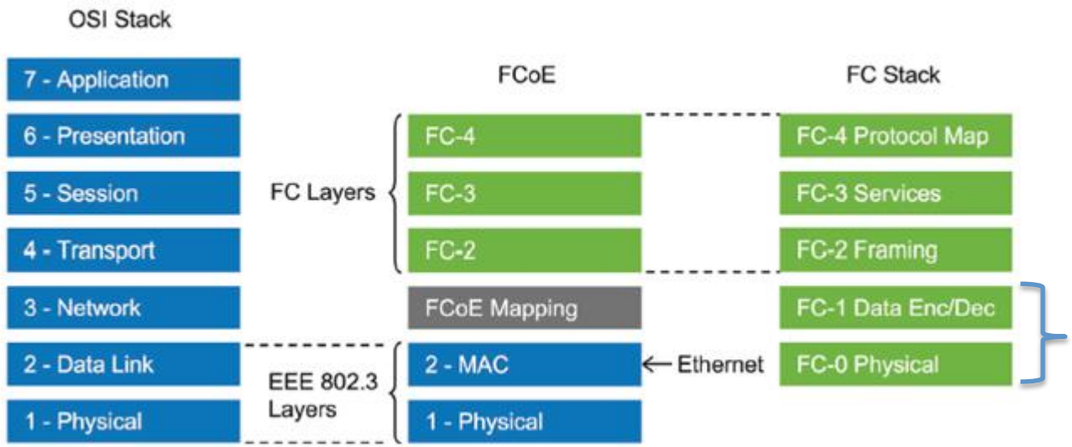

FC has functional layers:

- **FC-4**: Protocol-mapping layer, in which upper level protocols such as NVMe, SCSI, IP or FICON, are encapsulated into Information Units (IUs) for delivery to FC-2.
- **FC-3**: Common services layer, a thin layer that could eventually implement functions like encryption or RAID redundancy algorithms; multiport connections
- **FC-2**: Network Layer consists of the low level Fibre Channel protocols; port to port connections.
- **FC-1: Transmission protocol or data-link layer, encodes and decodes signals**
- **FC-0: The interface to the physical media; transceivers, cables, etc.**

Focus for this presentation

FC-PI-7 Starting requirements

- 64GFC had to be backward compatible to 32GFC and 16GFC
  - Backward compatibility and “plug and play” to utilize existing infrastructure with new speeds is always a must have for FC development
- Existing cable assemblies must plug into 64GFC capable products
  - LC (connector) and SFP+ (form factor)
- Reach goals
  - 100 meters for multimode short reach optical variant using OM4/OM5 cable plants
    - OM4 optical fibre has a higher optical bandwidth than OM3 fibre which leads to longer reach at a given speed
  - 10KM for single-mode optical variant
  - Electrical variant for backplane applications
- 64GFC doubles the throughput of 32GFC
- Corrected bit-error-rate (BER) target of 1e-15

Modulation changes for 64GFC

- Modulation refers to the signal levels that are on the “wire” (physical interface) whether optical or electrical
- Fibre Channel uses non return-to-zero (NRZ) or PAM2 modulation for 32GFC and slower
- 32GFC has a NRZ/PAM2 line rate of 28.05Gb.
- For 64GFC studies by the Fibre Channel committee and other committees determined that moving to PAM4 modulation would be “easier” from a component and IP perspective than staying with the NRZ/PAM2 modulation and doubling the “wire” rate to 57.8Gb
  - Actual “wire” rate for 64GFC is slightly more than double 32GFC
    - This will be explained a little later in this presentation

Signaling rate abbreviations

Abbreviation	Signaling rate	Number of Lanes		Data rate
1GFC	1.0625 GBd	1		100 MB/s
2GFC	2.125 GBd	1		200 MB/s
4GFC	4.250 GBd	1	8B/10B encoding	400 MB/s
8GFC	8.500 GBd	1	64B/66B encoding	800 MB/s
16GFC	14.025 GBd	1	256B/257B encoding	1600 MB/s
32GFC	28.050 GBd	1	256B/257B encoding	3200 MB/s
64GFC	28.900 GBd	1		6400 MB/s
128GFC	112.200 GBd	4		12800 MB/s
256GFC	115.600 GBd	4		25600 MB/s

MB/s = Mega*bytes* per second
GBd = Giga*baud* per second

Modulation changes for 64GFC (cont.)

- It was agreed to that PAM-4 modulation be used and the actual PAM4 wire rate is 28.9Gb
  - Two bits per symbol (one unit interval has two logical values) yield the needed 57.8Gbps
    - Throughput 6400MB/sec
  - Can run NRZ/PAM4 at a rate of 57.8Gbps or increase the number of bits per symbol to 2 bits per symbol in PAM4 and still get 57.8Gbps but the baud rate on the “wire” was 28.9Gb
    - Trade off between smaller symbol widths in the time domain versus smaller symbol heights in the voltage domain
- PAM4 modulation also uses Gray coding.
  - Gray coding implies that only one bit changes per step, i.e. 00, 01, 10, 11 would be the Gray coded PAM4
  - If a particular level is mis-interpreted, only one bit has transitioned so you don’t have two bits in error

Each PAM4 signal level corresponds to a two-bit symbol

PAM 2-NRZ

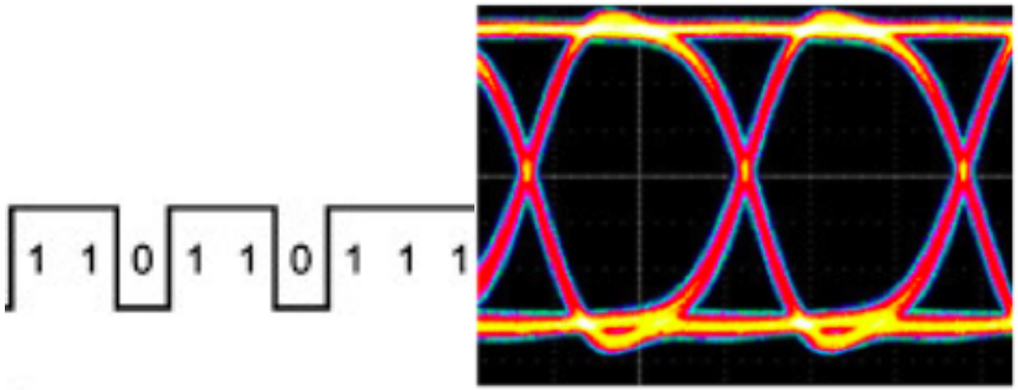

PAM 4

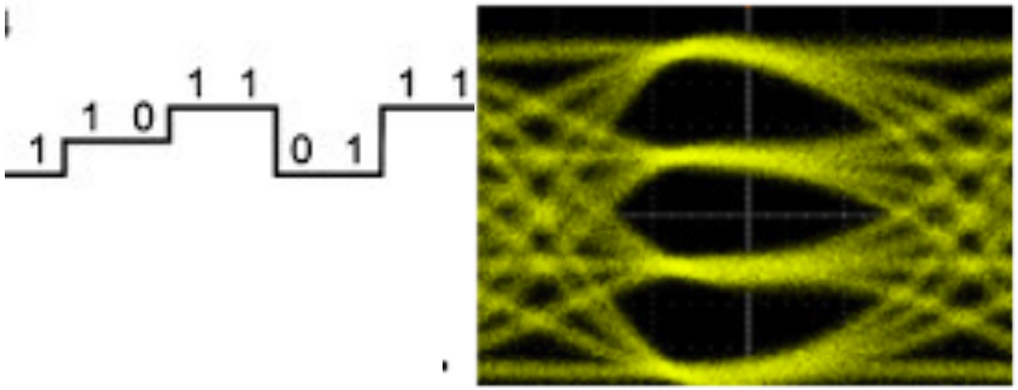

As serial data rates surpass 32Gb/s per channel, signal impairments necessitate the high-speed serial data technology to shift from simple NRZ (non-return to zero PAM2) signal modulation to the bandwidth efficient PAM4 (4-level pulse amplitude modulation)

Fibre Channel Variants in FC-PI-7

SM OS2	64GFC-LW	1300nm	0.5 m-**10km**	sub-clause 5.4
MM 50m OM3	64GFC-SW	850nm	0.5 m-70m	sub-clause 5.5
MM 50m OM4,OM5	64GFC-SW	850nm	0.5 m-**100m**	sub-clause 5.5
Backplane	64GFC-EA			clause 7

- Long Distance FC Webcast: https://www.brighttalk.com/webcast/14967/277327




OM3/OM4/OM5 considerations

OM3/OM4/OM5 - are incorporated into the FC-PI-7 standard

- OM5 offers improved distance capabilities over OM3 and OM4, for anticipated future WDM transmission
- Effective modal bandwidth for OM5 is specified at both 850 nm and 953 nm
  - ✓ OM4 and OM5 have the same 850nm EMB 4700MHz.Km
  - ✓ work is going on at IEC for OM3 and OM4 EMB at longer wavelengths up to 953 nm
  - ✓ Lime (lime green) is the official OM5 jacket color
  - ✓ Aqua is the official OM4/OM3 jacket color

- OM5/OM4 supports 64/32/16GFC at 100m distances at 850nm (16GFC supports 125m distance)
- ➢ FC Cabling Webcast: https://www.brighttalk.com/webcast/14967/303881
- The difference between OM3 (70m) 2000MHz.Km and OM4 (100m) 4700MHz.Km fiber is effective modal bandwidth at 850nm
  - this translates to longer FC frame transmission distances for the OM4
  - OM4 provides the most economical value proposition for single lane 64GFC
  - OM3/OM4 supports multiple wavelengths but at a shorter distance than is provided with OM5

MPO/MTP - Multiple-Fiber Push-On/Pull-off
LC - Lucent Connector

Fibre Channel Link

- The diagram below represents a simple Fibre Channel link.
- The link budget analysis looks at all the electrical and optical impairments end to end to determine if a transmitted signal can be received with a bit error rate below the minimum required.

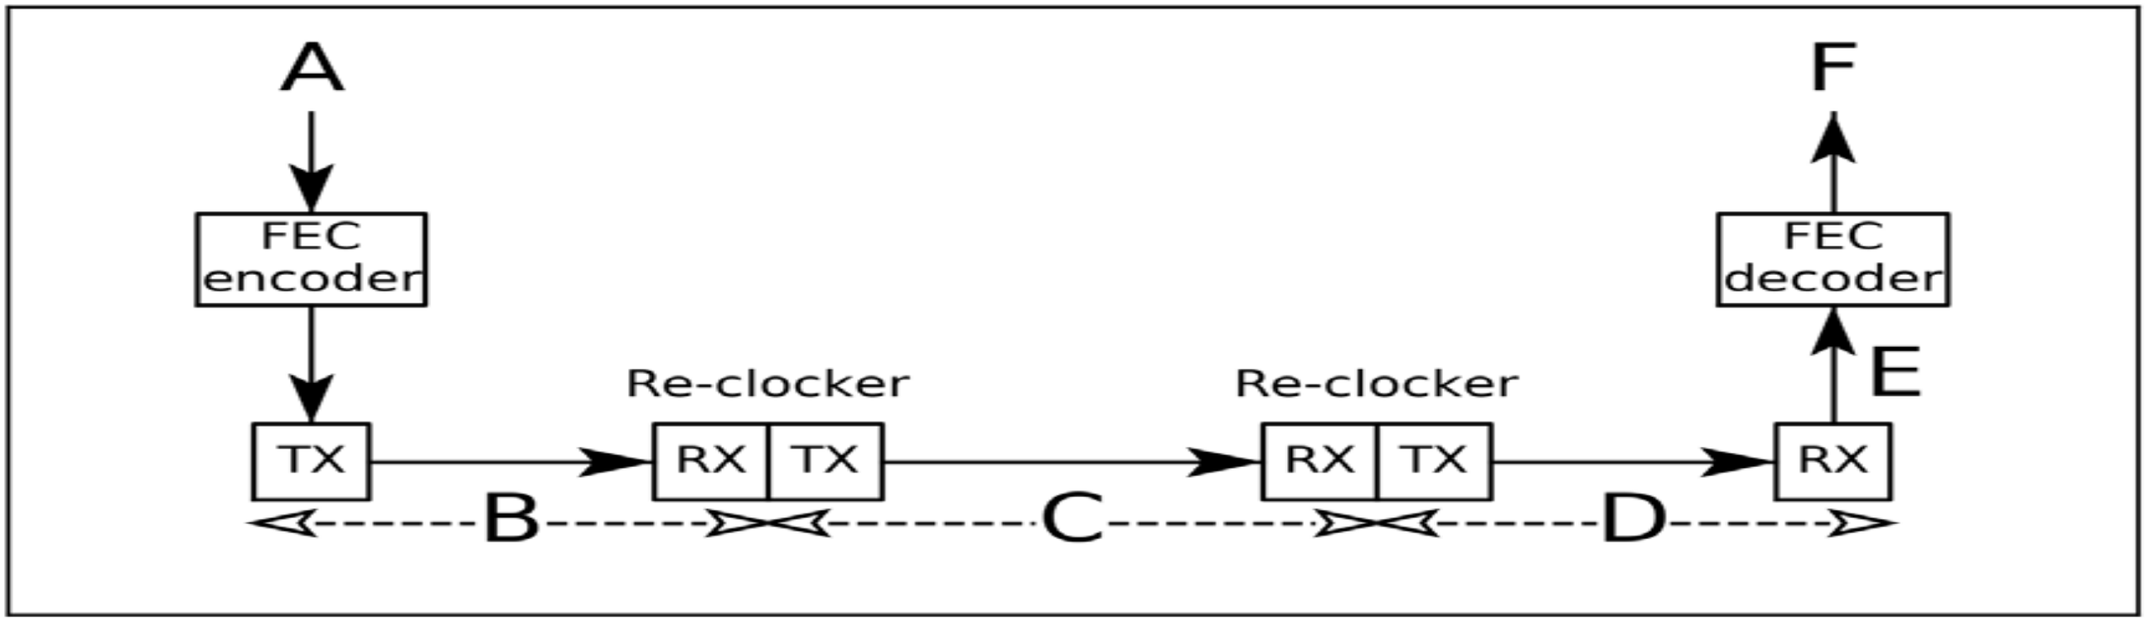

Optical transceivers for 64GFC

- Optical transceivers used for both short reach (multimode) and long reach (single-mode) utilize PAM4 signaling on the optical cable
- Electrical signal presented to the optical transceiver is PAM4 encoded
  - Optical transceiver transmits on the optical cable using PAM4
    - no modulation conversion is needed in the optical transceiver
- The optical transceiver has a clock and data recovery circuit (CDR) in the module on both the transmit and receive path
  - CDR resets the jitter budget at each optical transceiver
    - needed to close the link budget for the end to end link

64GFC SR and LR SFP MSA Form Factor

- 64GFC PAM4 electrical I/O, PAM4 optical I/O
- Tri-rate 64GFC/32GFC/16GFC SFP56 capability
- Retimer Technology
- Standard Diagnostics per SFF-8472

64GFC SW and LW
SFP28/SFP56 Form Factor

- SFP28 are speed enhanced variations of the predecessor SFP+ form factor, to support 32GFC
- SFP28 electrical and mechanical specifications were developed under SFF Committee
  - ✓ document SFF-8402 as well as earlier SFF-8472 and SFF-8432 specifications
- SFP56 is a general name for a 56G capable transceiver

64GFC Transmit Training

- Transmit Training Frame
  - Used for 32GFC and 64GFC link speed negotiation (LSN)
- 64GFC Transmit Training Frames
  - Gives the receive equalizer time to adapt to the PAM4 signal after speed negotiation
  - Sent, although the optical links don't run transmit training
  - Status frame indicates that fixed coefficients are being used, so training the link partner transmitter coefficients isn't needed
- Use in Speed Negotiation
  - Runs at a default PAM2 line rate of 28.05Gb
  - The control frame will advertise if PAM4 encoding is being used
    - Could also be PAM4 encoding with precoding enabled
  - If PAM4 encoding is agreed to between link partners then the line rate will be changed to PAM4 at 28.9Gb at the completion of link speed negotiation

64GFC speed negotiation

- The control and status frames are exchanged between link partners using Differential Manchester Encoding (DME)
- The DME run at 1/8 of the nominal signaling rate
- The DME coding runs at a lower signaling rate so that there is a high probability that the training control and status frames can be exchanged error free
- This was key for passive copper variants because the link partner information was exchanged before transmit training optimized FIR coefficients
- Transmit training in optical is run with fixed transmitter coefficients as advertised in the control frame

Forward error correction for 64GFC

- Forward Error Correction (FEC) is mandatory for all types of 64GFC links
- How it works
  - The transmitter encodes the data stream in a redundant way using an error correcting code
- 64GFC uses a block code called Reed Solomon.
  - The particular code used for 64GFC is RS(544,514)
  - This particular code allows correction of single bit errors or burst errors for 15 ten bit symbols out of 5140 bits sent
- 64GFC uses terms such as uncorrected BER which is the minimum BER to be expected pre-FEC encoding/decoding
  - Uncorrected BER is in the 1e-04 range or lower for 64GFC
  - FEC corrected BER is in the 1e-15 range or lower for 64GFC
  - These numbers help identify the usefulness of FEC in making 64GFC links robust

Forward Error Correction (FEC)

A set of algorithms that perform corrections that allow for recovery of one or more bit errors

- SNIA Dictionary

Forward error correction for 64GFC

- FEC has been used in previous FC variants
  - 32GFC had mandatory RS(528,514) FEC
  - 16GFC had optional (2112,2080) FEC which was primarily used for copper variants
- The small baud rate increase for 64GFC (28.9Gb PAM4) versus 32GFC (28.05Gb PAM2) can be explained by the more powerful FEC used for 64GFC
- The extra parity bits needed for the 32GFC RS(528,514) came from transcoding the 64/66 bit stream using a 256/257 transcoder. This allowed gaining 140 parity bits without increasing the baud rate
- For 64GFC RS(544,514) is used and an additional 160 parity bits are needed and the only way to get the parity bits is to run at a slightly higher baud rate of 3.03%, i.e. 28.9Gb versus 28.05Gb

256GFC (parallel four lane)

- FC-PI-7P will describe a four lane 64GFC variant that has a throughput off 256GFC
- The FC-PI-7P standard is currently in development
  - Expect the first letter ballot to be in the middle of 2019
- The data is stripped across the four lanes
- MRD requested the following variants
  - 100 meters on multimode cable plants OM4/OM5
  - 2km single-mode variant
- Backward compatibility with 128GFC is also a requirement (4x32GFC)

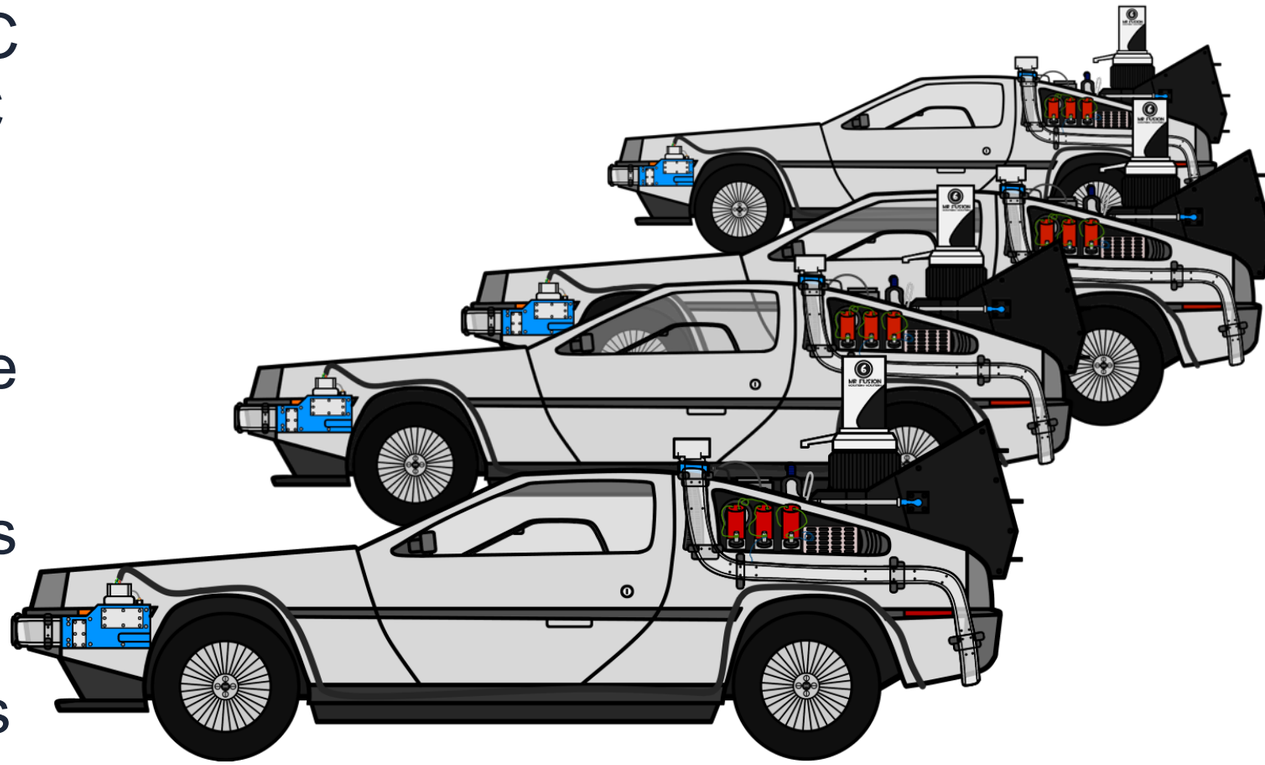

128GFC FC-PI-8 Planned Requirements

- Backward compatible to 64GFC and 32GFC
- Same external connectors as 32/64GFC
- Existing cable assemblies will work with 128GFC
- Multimode cable plant reach is 100 meters on OM4/OM5
- Single-mode cable plant reach of 10KM
- 128GFC links should double the throughput in MB/sec of 64GFC links
- Corrected BER target of 1e-15
- Reduce latency of 64GFC by up to 20%

Summary

- FCIA publically announced the completion of the FC-PI-7 standard enhancements, signaling the industry to begin development
- The fibre channel industry has quickly implemented the FC-PI-7 standards and physical layer and protocol testing was initiated
- 64GFC products are ready to enter the FC SAN market and take fibre channel performance to the next level
- Parallel 64GFC is being specified

✓ Who could ask for more

After this Webcast

- Please rate this event – we value your feedback!

- We will post a Q&A blog at http://fibrechannel.org/ with answers to questions received today
- Follow us on Twitter @FCIAnews for updates on future FCIA webcasts

- Visit our FCIA webcast library at http://fibrechannel.org/webcasts/ to learn about:
  - Fibre Channel Fundamentals https://www.brighttalk.com/webcast/14967/255009
  - FC-NVMe https://www.brighttalk.com/webcast/8609/242341
  - Long Distance Fibre Channel https://www.brighttalk.com/webcast/14967/277327
  - How to Use Fibre Channel Speed Map https://www.brighttalk.com/webcast/14967/246353
  - FCIP (Extension): Data Protection and Business Continuity https://www.brighttalk.com/webcast/14967/285643
  - Fibre Channel Performance https://www.brighttalk.com/webcast/679/295141
  - Fibre Channel Cabling https://www.brighttalk.com/webcast/14967/303881
  - FICON 101 https://www.brighttalk.com/webcast/8609/315723
  - Fibre Channel Interoperability https://www.brighttalk.com/webcast/8615/327269
  - Protocol Analysis for High-Speed FC Fabrics https://www.brighttalk.com/webcast/8615/333863

Our Next FCIA Webcast:

FICON 201 – February 20th 10:00 am PT

Register at:

https://www.brighttalk.com/webcast/14967/344257

Thank You