

Physical Layer Tests of 100 Gb/s Communications Systems

Application Note

Table of Contents

1. Introduction	3	4. Diagnostic tests	23
2. 100G and related standards	4	4.1. What to do if the transmitter fails.....	23
2.1. 100 GbE – IEEE Standards 802.3ba, 802.3bj, and 802.3bm.....	6	4.2. What to do if the receiver fails	24
2.2. OIF-CEI.....	6	Summary.....	24
2.3. Fibre Channel 32GFC.....	6		
3. Testing 100G Systems.....	7		
3.1. Testing optical transmitters.....	9		
3.2. Testing optical receivers	11		
3.3. Testing electrical transmitters	16		
3.4. Testing electrical receivers.....	19		

1. Introduction

The list of bandwidth intensive applications has exploded in the last decade. Video-on-demand, voice-over-IP, cloud-based computing, and storage have created a ravenous bandwidth appetite that has rushed deployment of 100 Gb/s technology and spurs development of 400G.

The power of HSS (high speed serial) technology, with its noise-resistant differential signaling and jitter-resistant embedded clocking plus closed-eye equalization, enables 25+ Gb/s on previously inconceivable lengths of PCB (printed circuit board). Combining HSS links in parallel simplifies electrical 100G signal transmission to optical transceivers, easing connectivity to the fiber optic backbone. The result is that datacom and telecom technologies use 100 Gigabit Ethernet (100 GbE) for transport, including SAS, Infiniband, even Fibre Channel.

This note covers the transmitter and receiver tests necessary to assemble 25-32 Gb/s single lane systems and complete 100G systems. Since every 25+ Gb/s HSS technology shares common themes, we'll follow 100 GbE, that is, the IEEE 802.3ba, 802.3bm, and 802.3bj compliance specifications but point out differences between other high rate systems like Fibre Channel's 32GFC, and the Optical Internetworking Forum's CEI (common electrical interface) implementation agreements, called IAs rather than standards.

As we work through the tests, we'll encounter common themes in the interplay of jitter, noise, frequency response, and crosstalk. After the description of compliance tests, we'll suggest strategies for diagnosing noncompliant components and systems as well as techniques for measuring performance margins.

Technology specifications tend to be written in the engineering equivalent of legalese, so we composed this note as a supplement to clarify the tests themselves, their role, and how to perform them.

Standard		Geometry	Reach or Max Loss	Data Rate
100 GbE	100GBASE-ER4 100GBASE-LR4	4 SM fibers	2 m to 30 km 2 m to 10 km	4×25.78125 Gb/s
	100GBASE-SR4	4 MM fibers	≥ 100 m	
	100GBASE-CR4	4 sets of matched cables	≥ 5 m	
	100GBASE-KR4	PCB	≤ 35 dB at 12.9 GHz	
OIF-CEI	CEI-25G-LR	PCB	0-68.6 cm + 1 or 2 connectors	19.90-25.80 Gb/s
	CEI-28G-MR	PCB	0-50 cm + 1 connector	19.90-28.10 Gb/s
	CEI-28G-SR	PCB	0-30 cm	19.90-28.05 Gb/s
	CEI-28G-VSR	PCB	10 cm on host PCB + 1 connector + 5 cm on module PCB	19.60-28.10 Gb/s
Fibre Channel	32GFC	N channels optical and electrical	Many optical and electrical variations	28.05 Gb/s

Table 1. Summary of 100G and related standards.

2. 100G and related standards

Standards recommend tests to assure component interoperability. In this section, we summarize the specifications, Table 1. It's important to keep in mind that technology standards documents constantly evolve. The numbers we quote should be considered typical of what to expect, but for compliance testing, always check the actual standards!

At the electrical end, the technology shares the following characteristics: balanced, unidirectional, 100 Ohm differential signaling with embedded clocks, low voltage swings, NRZ (non-return to zero) signals, and multiple channels. As data rates increase and engineering experience grows, multi-level signaling, particularly PAM4 (4-level pulse amplitude modulation), will replace NRZ in certain applications. By encoding two bits in each symbol, PAM4 reduces the bandwidth demand at a given data rate by about half, at the expense of signaling complexities. This note focuses on NRZ technology. PAM4 is not addressed here.




Figure 1. Reed Solomon forward error correction scheme. RS-FEC(n, k) is capable of correcting up to t errors in each word.

FEC (forward error correction) is required in some of the 100G specifications and optional in others. Reed-Solomon FEC is like a cyclic redundancy check but with the ability to correct as well as count errors. The RS-FEC(n, k) scheme uses an algebraic technique to encode k symbols of data with $2t$ parity symbols into an n symbol word in such a way that t symbol errors can be corrected, Figure 1.

For example, all of the electrical (100GBASE-CR4 and KR4) and one of the optical (100GBASE-SR4) 100 GbE specifications covered in this document use RS-FEC(528, 514). The PAM4 100GBASE specifications uses a different RS-FEC scheme.

Here's how it works: each symbol consists of ten bits. Each codeword consists of 528 symbols, totaling 5280 bits. RS-FEC(528, 514) can correct at most 7 symbols. The ability of RS-FEC(528,514) to *identify* errors isn't so restricted.

The capacity for FEC to correct errors relaxes the BER compliance requirement at the expense of latency. To account for FEC limitations, both the BER and FLR (frame loss rate) are specified. Where BER is defined as the ratio of the number of bits in error to the total number of bits transmitted, the FLR is defined as the ratio of the number of invalid, lost frames, to the total number of frames transmitted.

To meet the operational net BER $< 10^{-12}$ requirement, 100 GbE specifies both:

BER $< 5 \times 10^{-5}$ prior to or independent of FEC and

FLR $< 6.2 \times 10^{-10}$ for 64 octet (i.e., 512 bit) frames, including FEC.

Since data from each channel is striped prior to FEC operation, these requirements nearly assure operation at BER $< 10^{-15}$ even in the presence of burst errors.

The two electrical applications for which FEC is not required, 100GBASE-ER4 and LR4, require BER $< 10^{-12}$.

Since terminology varies among the specifications, let's clear up potential misunderstandings immediately. In this document, we distinguish the data rate and payload rate; the data rate is the rate at which raw data propagates. The payload rate does not include overhead from FEC and coding, hence, payload rate $<$ data rate. Since we'll only discuss NRZ signaling, we use Gb/s rather than Gbaud and restrict the term "symbol" to RS-FEC's 10 bit symbol fields.

2.1. 100 GbE – IEEE Standards 802.3ba, 802.3bj, and 802.3bm

Three IEEE standards cover the 100 GbE systems listed in Table 1. All of them achieve the nominal 100 Gb/s data rate by combining four separate 25.78125 Gb/s lanes.

The IEEE Std 802.3ba standard covers the long reach, 100GBASE-LR4 and extended reach, 100GBASE-ER4, four lane single mode fiber specifications over 10 and 40 km, respectively. The differences between them are primarily at the receive end. The ER4 receiver has greater sensitivity and has to pass a more difficult stress tolerance test than the LR4 receiver.

The IEEE Std 802.3bm covers the short reach 100GBASE-SR4 MM (multimode) fiber specification which has a range of at least 100 m. Where RS-FEC is optional for 100GBASE-LR4 and 100GBASE-ER4, it's mandatory 100GBASE-SR4.

The IEEE std 802.3bj covers the electrical specifications. 100GBASE-CR4 consists of four lanes of shielded balanced cables with reach up to 5 m and 100GBASE-KR4 consists of four differential traces on backplanes. Rather than specify a minimum reach, KR4 specifies the maximum allowed insertion loss at the Nyquist frequency which provides greater system design flexibility. RS-FEC is mandatory for both 100GBASE-CR4 and 100GBASE-KR4.

2.2. OIF-CEI

Implementation Agreements (IAs) from OIF-CEI do not prescribe compliance tests the way that IEEE's 802.3 100 GbE or Fiber Channel specifications do. Instead, the emphasis is on informative and normative tests that attempt to assure component interoperability across standards. "Normative" tests are like compliance tests in the sense that the committee prescribes them to assure interoperability. "Informative" tests are recommended to develop a more thorough understanding of performance and margin.

In this note, we draw from the OIF-CEI-03.1 IAs summarized in Table 1.

The long reach CEI-25G-LR consists of multiple lanes at 19.90-25.8 Gb/s of differential pairs over up to 686 mm of PCB trace with up to two connectors. While the PCB channel is specified to a length up to 686 mm, the physical length is less important than the channel frequency response. The medium reach CEI-28G-MR consists of multiple lanes at 19.90-28.10 Gb/s for signaling of differential pairs over up to 500 mm of PCB trace, vias, and up to one connector. The short reach CEI-28G-SR consists of multiple lanes at 19.90-28.05 Gb/s of differential pairs over 300 mm of PCB trace, vias, and up to one connector.

The very short reach CEI-28G-VSR consists of multiple electrical lanes at 19.60-28.05 Gb/s for signaling between serdes (called hosts in the IA) and transceivers (modules in the IA). The serdes and transceiver can be separated by some 100 mm of host PCB trace to a mated connector pair plus an additional 50 mm or so of transceiver trace.

The OIF-CEI provides independent IAs for each level of the protocol stack in order to permit mixing and matching of applications, like use of a SONET framer in an optical module connected to a CEI electrical interface. Since FEC operates just above the media level, FEC is optional in OIF-CEI.

2.3. Fibre Channel 32GFC

The high rate Fibre Channel standard, 32GFC, has a data rate of 28.05 Gb/s. The confusing name scheme, 32GFC for 28.05 Gb/s technology, comes from the desire for the name of each generation to demonstrate that the payload rate, as opposed to the data rate, is double that of the previous generation.

The confusion began with a large decrease in overhead in the transition from 8GFC to 16GFC when the data rate advanced from 8.5 to 14.025 Gb/s but the payload rate doubled from 6.4 to 12.8 Gb/s. The payload rate for 32GFC is 25.6 Gb/s, twice that of 16GFC, but the data rate, 28.05 Gb/s, is well short of that implied by the 32GFC abbreviation.




Figure 2. Diagram of (a) 4x25 Gb/s 100G serdes-transceiver WDM optical system, (b) 4x25 Gb/s 100G serdes-transceiver optical system, and (c) a 4x25 Gb/s 100G serdes to serdes electrical system. The figure does not show the symmetric return paths.

3. Testing 100G Systems

Figure 2 is a diagram of the components of typical 100G systems. A serdes serializes a signal and transmits four 25+ Gb/s differential pairs. The serdes may consist of either separate components for each output or a single integrated component. The 25+ Gb/s electrical signals are transmitted from the serdes to an optical transceiver. The transceiver retimes the signals and transmits optical versions on either SM (single mode) or MM (multimode) optic fibers. A second transceiver receives the optical signals, converts them to electrical signals and transmits them to another serdes for deserializing. The purely electrical version follows the same scheme without the intermediate transceiver-driven optical signaling.

Whether for transmitter or receiver testing, optical or electrical, we need test patterns that put every aspect of a component and every component of a system to the test. The PRBS n (pseudo-random binary sequences of length 2^n-1) are standardized patterns with every permutation of n bits. The CEI CID jitter tolerance pattern is designed to have the most aggressive elements of the PRBS31, plus 72 CID (consecutive identical) bit sequences but at a more manageable total length than PRBS31.

Test Patterns	
0x00ff square wave	8 bits low, 8 bits high
PRBS9	511 bits
PRBS15	32,767 bits
PRBS31	2.1 Gbits
RS-FEC encoded scrambled idle	
CEI CID jitter tolerance pattern	(72 CID bits + ≥ 10328 from PRBS31 + seed) + complement

Table 2. Test Patterns.

All PatternPro and BERTscope pattern generators provide all test patterns used in 100G communications, including PRBS31, RS-FEC scrambled idle, or for that matter, every common test pattern as well as any that you devise up to 2 MB long.

All transmitter tests, both electrical and optical, should be performed with every system channel active in both directions to include all reasonable sources of crosstalk interference. To prevent unrealistic data-dependent interference, patterns on the crosstalk channels should differ from the test signal pattern. If it's not possible for each aggressor to transmit a unique pattern, introduce sufficient delay between them so that the patterns aren't synchronized.

Since the frequency response of PCB punishes high frequency content, electrical signaling between serdes chips or between serdes and transceivers across centimeters of PCB requires signal conditioning: pre-emphasis at the transmitter and equalization at the receiver.

The big difference between the 100 GbE electrical specifications and previous HSS data specifications is extensive characterization of channels in terms of the frequency dependence of insertion and return loss and the introduction of COM (channel operating margin). Since frequency response determines ISI (inter-symbol interference), compliant channels restrict ISI to levels that can be accommodated by the combination of transmitter equalization—usually three taps of pre- or de-emphasis, a

form of FFE (feed-forward equalization) as it is referred to in some of the standards documents—and receiver equalization that includes a CTLE (continuous time linear equalization) filter usually combined with a DFE (decision feedback equalizer). Since 100G technology requires both transmitter and receiver equalization, channel insertion loss must meet both minimum and maximum criteria.

Other differences include more careful parameter definition like UUGJ (unbounded uncorrelated Gaussian jitter) in place of RJ (random jitter) and introduction of UBHPJ (unbounded high probability jitter) which combines PJ (periodic jitter) and crosstalk. EBUJ (effective bounded uncorrelated jitter) and ETUJ (effective total uncorrelated jitter) are terms derived from the dual-Dirac model applied to the components of the jitter distribution that are uncorrelated to the data.

COM (channel operating margin) has been introduced as a way to combine the signal impairments from the transmitter and channel into one signal-to-noise-like parameter. COM is derived by measuring channel S-parameters and then modeling the effect of crosstalk from other channels, and random noise and jitter, to calculate the ratio in dB of signal amplitude to aggregate noise amplitude at the target detector error ratio.

By specifying a maximum compliant COM, the standards grant designers latitude in budgeting the combination of ISI, random noise, random jitter, and crosstalk signal impairments as they see fit in a way that still assures interoperability.

Summary of Typical 100 GBE Optical Transmitter Requirements

	100GBASE-ER4	100GBASE-LR4	100GBASE-SR4
Average launch power	-2.9 to 2.9 dBm	-4.3 to 4.5 dBm	-9 to 2.4 dBm
OMA	-1.3 to 4.5 dBm	0.1 to 4.5 dBm	-7 to 3 dBm
Extinction ratio	≥ 8 dB	≥ 4 dB	≥ 2 dB
Eye mask {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.4, 0.45, 0.25, 0.28, 0.4}		{0.3, 0.38, 0.45, 0.35, 0.41, 0.5}

Table 3. Summary of transmitter specifications.


3.1. Testing optical transmitters

Typical transmitter requirements are summarized in Table 3.

Figure 3 defines the eye diagram, average launch power, OMA (optical modulation amplitude), and, together with the values in Table 3, the eye mask criteria. The normalized logic 0 and 1 levels used in the eye-mask are defined by the averages of the lower and upper halves of the central 0.2 UI of the eye. Eye masks are the do-not-enter regions of the eye diagram specified by six parameters, as shown in both Table 3 and Figure 3.

Eye mask tests can be performed on a DSA8300 low-noise equivalent-time sampling oscilloscope or BERTscope. In both cases, wide bandwidth optical-to-electrical receivers and clock recovery units are necessary. The 3 dB clock recovery bandwidths differ among specifications and can be accommodated by CR286A, a digital-based, second order Phase-Locked Loop (PLL) with user-specified corner frequencies capable of tracking jitter to 23 MHz. The clock recovery bandwidth for optical 100 GbE systems has a corner frequency of 10 MHz with slope of 20 dB/decade.

Apply crosstalk by engaging pattern generators on three other channels or however many parallel lanes your system supports. Each crosstalk signal should meet the requirements

**Figure 3.** Definitions of the transmitter eye mask and OMA.

for a compliant transmitter and their test patterns should differ from the test signal pattern. If it's not possible for each aggressor to transmit a unique pattern, introduce sufficient delay between them so that the patterns aren't synchronized.

The optical-to-electrical receiver should apply a 4th order Bessel-Thompson filter with a reference frequency of three-fourths the data rate, $\frac{3}{4} f_{\text{data}}$. The filter is required so that different test platforms can operate under uniform measurement conditions.




Figure 4. Eye mask measurement of a 100 GbE optical 25 Gb/s signal.

The random nature of a mask test is addressed by requiring a minimum “hit ratio.” The hit ratio is defined as the ratio of the number of mask violations to the total number of samples acquired per unit interval. Since this is a statistical measurement, it’s worthwhile to keep in mind that accuracy improves with more hits.

A transmitter is compliant if it achieves a **hit ratio less than 5×10^{-5}** , Figure 4.




Figure 5. Eye mask testing with BER contours. The BER= 10^{-6} contour, the outer blue contour, corresponds to a 5×10^{-5} hit ratio.

Alternatively, it’s both more statistically reliable and easier to measure the BER Contour on a PatternPro Error Detector, BERTscope, or on a DSA8300 oscilloscope equipped with 80SJNB jitter and noise analysis software. As long as the **BER= 10^{-6} contour is outside the mask**, Figure 5, the transmitter passes the 5×10^{-5} hit ratio eye test. This technique also makes it easier to see the margin with which a transmitter passes.

Summary Of 100 GbE Optical Receiver Test Conditions		
Long and Extended Reach	100GBASE-ER4	100GBASE-LR4
Average received power	4.5 to -20.9 dBm	4.5 to -10.6 dBm
Optical Modulation Amplitude (OMA)	≤ -17.9 dBm	≤ -6.8 dBm
Conditions for stressed receiver tolerance testing:		
Vertical Eye Closure Penalty (VECP)	3.5 dB	1.8 dB
SJ	According to the template in Figure 7	
Sinusoidal interference 0.1-2 GHz	Sum to J2 and J9	
RJ		
J2 jitter	0.3 UI	
J9 jitter	0.47 UI	
BER requirement	$\leq 10^{-12}$	
Short Reach		
100GBASE-SR4		
Average received power	-10.9 to 2.4 dBm	
Optical Modulation Amplitude (OMA)	≤ 3 dBm	
Stressed receiver sensitivity	-5.6 dBm	
Conditions for stressed receiver sensitivity testing:		
Extinction ratio	≥ 2 dB	
Stressed eye closure (SEC)	4.9 dB	
SJ	According to the template in Figure 7	
Sinusoidal interference	Adjust RJ, SI, and extinction ratio to meet the SEC, J2, J4, and mask requirements.	
RJ		
J2 jitter	0.39 UI	
J4 jitter	0.53 UI	
Aggressor OMAs	3 dBm	
Stressed mask	{0.28, 0.5, 0.5, 0.33, 0.33, 0.4}	
BER for any SJ amplitude and frequency in Figure 7	$\leq 10^{-12}$ or BER prior to FEC $\leq 5 \times 10^{-5}$ or FLR (frame loss ratio) post FEC $\leq 6.2 \times 10^{-10}$	
BER without FEC		
or BER prior to FEC		
or FLR (frame loss ratio) post FEC		

Table 4. Summary of ER, LR, and SR 100 GbE stressed receiver sensitivity test conditions.

3.2. Testing optical receivers

The 100 GbE optical receiver stress tolerance tests for the extended and long reach (100GBASE-ER4 and LR4) single-mode fiber technologies are quite similar except that greater

sensitivity and robustness is required for extended reaches, as indicated in Table 4. The short reach (100GBASE-SR4) multimode fiber specifications differ substantially from the ER4 and LR4 specifications so we've noted them separately in Table 4.

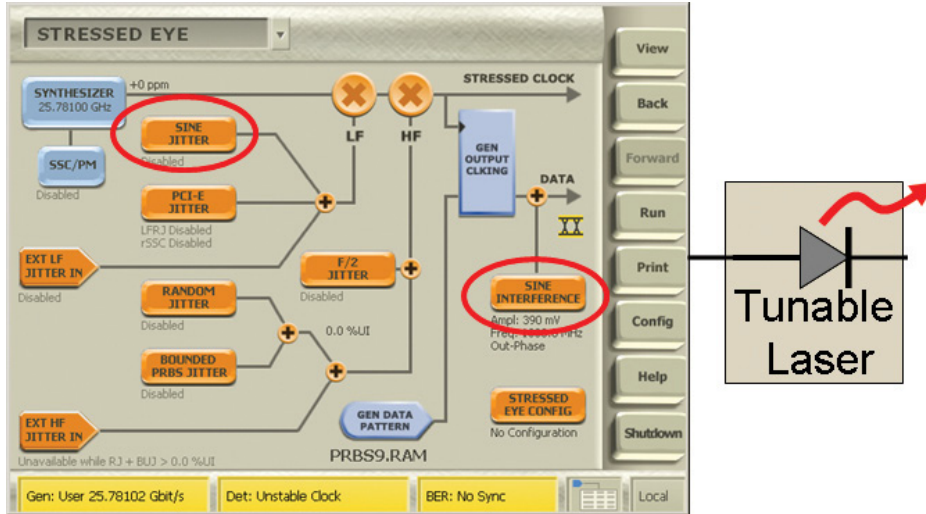


Figure 6. Configuring the stressed receiver test signal.

Figure 6 shows how to apply calibrated levels of stress to the test signal and Table 4 summarizes the stresses. At these data rates, producing compliant stress levels is tricky. Both the PatternPro and BERTScope Pattern Generators generate conformant, stressed eyes using their internal impairment systems by driving a tunable laser-based signal into the optical receiver device under test, Figure 6.

The receiver *tolerance* tests required of ER4 and LR4 are similar to SR4's receiver *sensitivity* test so we'll go through them together but note their differences.

First, **configure the pattern generator** to drive a Mach-Zehnder (MZ) optical modulator. Then tune the MZ bias to optimize 1/0 symmetry, but don't exceed the OMA given in Table 4. The pattern generator should transmit either a PRBS31 or RS-FEC scrambled idle.

Introduce a **4th order low pass Bessel-Thompson** filter with a 3 dB roll off at 19 GHz to the pattern generator output. This filter serves three purposes: first, it removes higher order harmonics from the pattern generator which allows more consistent VECP (vertical eye closure penalty) measurements; second, it applies ISI; and, third, when combined with the response of the receiver, should result in a frequency response with the appropriate level of **SEC** (stressed eye closure).

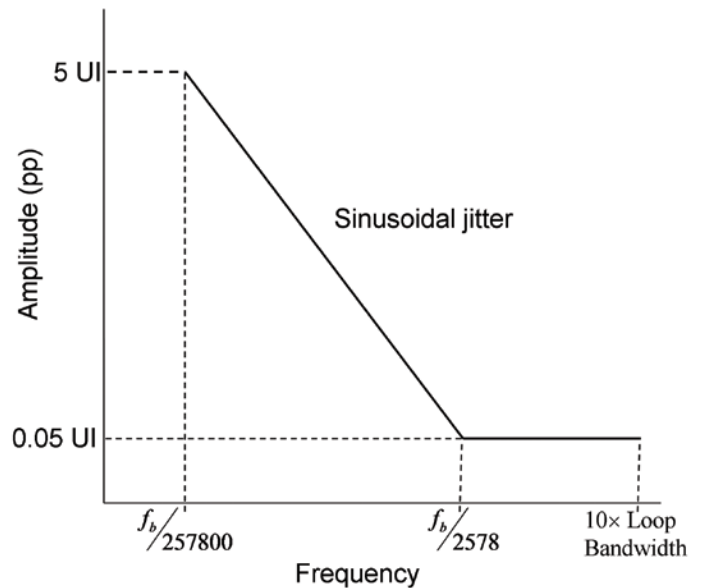



Figure 7. Sinusoidal Jitter stress template.

Apply the minimum value of **SJ** (sinusoidal jitter) to the pattern generator clock above the rolloff in the template shown in Figure 7 by using either a PatternPro Pattern Generator or a BERTScope.

SI (sinusoidal interference) amplitude modulation closes the eye both vertically and horizontally by shifting both the rails and non-zero slopes of logic transitions up and down which adds amplitude noise and converts amplitude noise to timing noise. Apply two sources of SI, one prior to and one following the limiting amplifier, as in Figure 6. Set their frequencies anywhere between 0.1 and 2 GHz, but make certain that the frequencies have no harmonic relationship with the SJ frequency, the data rate, the pattern repetition rate, or each other. A nice trick is to use prime numbers, say 13 MHz and 47 MHz. SI should not cause bit shrinkage larger than 0.1 UI.

Apply **UUGJ** (unbounded uncorrelated Gaussian jitter) by setting the **Random Jitter (RJ)** to the appropriate level with the precision Gaussian noise generator integrated in the PatternPro or BERTscope pattern generator.

Setting **VECP** to the level given in Table 4 for ER4 and LR4 (SR4 doesn't specify VECP) is a multi-step process. Optical VEC is given by:

$$VECP = 10 \log \frac{OMA}{EH(2.5 \times 10^{-3})}$$

where the eye height, $EH(2.5 \times 10^{-3})$, is the vertical eye opening defined at $BER=2.5 \times 10^{-3}$. While conceptually cumbersome, $EH(BER)$ is a more precise definition than average peak-to-peak voltage swing. It is equivalent to the vertical distance at the center of the eye between BER contours of 2.5×10^{-3} which is easy to measure with the PatternPro Error Detector, BERTscope, or DSA8300.

At this point SR4 differs from ER4 and LR4. All three require jitter stresses defined with respect to the overall jitter distribution. Due to the large statistical distributions required to probe BERs below 10^{-12} , the jitter distribution can be cumbersome. To make it more manageable without sacrificing rigor, constraints are placed on certain fractions of the jitter distribution.

J2, **J4**, and **J9** define fractions of the jitter distribution. J2, for example, is the time interval that contains all but 10^{-2} of the

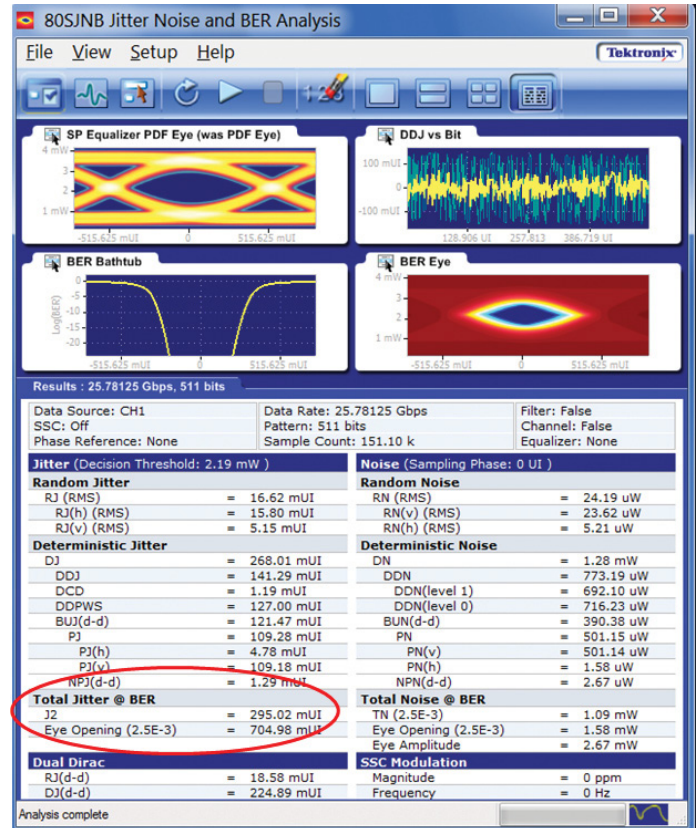


Figure 8. Example measurement of J2 on the DSA 8300 with 80SJNB.

distribution, which is where DJ is dominates. Hence J2 covers the jitter distribution from its center out to $BER = 2.5 \times 10^{-3}$. At the opposite extreme, J9 indicates the time interval of the outer billionth of the distribution, from $BER = 2.5 \times 10^{-10}$ and lower. J9 is dominated by the Gaussian tails of the jitter distribution. J4, on the other hand, is the time interval that includes all but 10^{-4} of the distribution, from its center out to $BER = 2.5 \times 10^{-5}$; where RJ and DJ are both prominent.

To assemble the necessary stress, the ER4 and LR4 standards specify J2 and J9 and SR4 specifies J2 and J4.

For ER4 and LR4, after setting the required VECP, add SI to the signal, until the J2 requirement is reached, Figure 8.

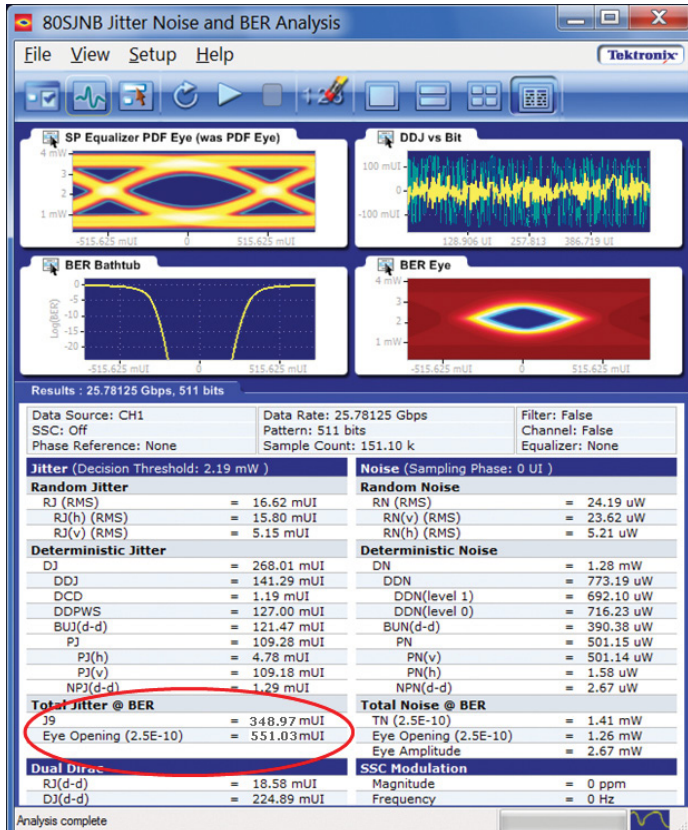


Figure 9. After setting J2, the J9 level is 0.35 UI, to meet the J9 spec, add less than 335 fs rms RJ.

Because of the wide BER disparity between J2 and J9, a tiny amount of rms RJ raises J9 to the required level with only a small effect on J2. For example, Figure 9, if after setting J2, the J9 level is 0.35 UI, add 0.12 UI of RJ to meet the 0.47 UI J9 spec—that’s less than 332 fs of rms RJ. Particular care needs to be taken to ensure the intrinsic RJ of the generator source is lower than 332 fs; otherwise meeting the J2 and J9

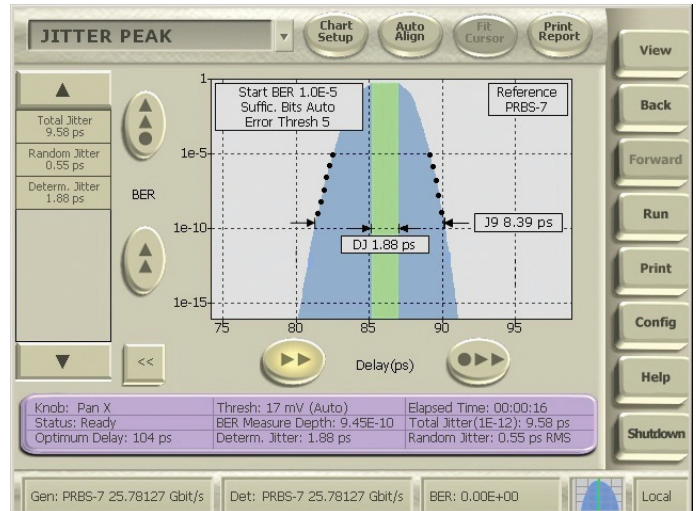


Figure 10. Jitter Peak with J9 measured on an error detector.

requirements simultaneously becomes impossible. Adding femtosecond levels of RJ above the instrumentation noise floor requires a precision RJ source. It’s vitally important to calibrate the system with the precise instrument coordination provided by combining either a PatternPro or BERTScope pattern generator and a DSA8300 oscilloscope with 80SJNB software, since it can derive J2 and J9 automatically.

Setting J4 for SR4 follows the same strategy as setting J2 and J9 for ER4 and LR4, except that the tradeoffs are somewhat easier to handle since J4 is less sensitive to tiny amounts of RJ: tune RJ, SI, and the transmitter extinction ratio until the **SEC**, J2, and J4 requirements in Table 4 are met simultaneously and the resulting stressed eye passes a standard eye-mask test with the mask defined in Table 4. The resulting extinction ratio should approximate the minimum specified.




Figure 11. Optical stressed receiver test.

Transmit the stressed signal into a replication of the system, **Figure 11**, along with three active aggressor signals for crosstalk. The aggressors don't have to meet the stress criteria.

If the receiver is capable of counting its own BER, FLR (frame loss ratio) after applying FEC, BER prior to FEC, or RS-FEC symbol error ratio, you're ready to go. If not, connect the receiver output to the BERTscope or PatternPro Error Detector. If the receiver doesn't provide a clock output you're supposed to time the error detector with a clock recovery unit, but since the receiver output has been retimed and should be pristine you can probably use the BERTscope or PatternPro Pattern Generator clock output instead.

Apply the stressed signal to the receiver first with low amplitude SJ applied above the roll off frequency, Figure 7.

If the receiver operates at or better than the specified error criterion, continue testing across the SJ frequency-amplitude template in Figure 7 to assure that the receiver can track low-frequency jitter with all other stresses applied. It's automatic with the PatternPro Error Detector and BERTscope jitter tolerance measurement functions.

For ER4 and LR4, the receiver must operate at $BER \leq 10^{-12}$ across the SJ template.

For SR4, the interface BER is given by the average BERs of the four receive lanes subjected to stress sequentially. Since SR4 uses RS-FEC, separate compliance criteria are given in Table 4 that depend on how the receiver counts errors; all of these criteria are equivalent in principal and ultimately assure net 100 Gb/s operation at $BER \leq 10^{-12}$.

Typical Electrical Transmitter Requirements	
Data rate (c.f., baud rate)	19.90-28.05 Gb/s
Peak-to-peak differential voltage	800-1200 mV 1200mV CR4/KR4
Rise/fall time (20/80%)	≥ 8-10 ps
Pre-emphasis	≥ 3 taps
Pulse response optimization	≥ 8 UI
Compliance test board with specified frequency response, $IL(f)$	-25 to -1.25 dB at Nyquist
CEI-28G-VSR: RMS common mode noise	≤ 17.5 mV
Even-Odd Jitter (EOJ, c.f., DCD)	≤ 0.035 UI
CR4/KR4: Effective bounded uncorrelated jitter (EBUJ)	≤ 0.1 UI
CR4/KR4: Effective total uncorrelated jitter (ETUJ) at BER=10 ⁻¹⁵	≤ 0.18 UI
CEI: Uncorrelated Bounded High Probability Jitter	≤ 0.15 UI
CEI: Uncorrelated Unbounded Gaussian Jitter at BER=10 ⁻¹⁵	≤ 0.15 UI
CEI-28G-VSR: Eye width at BER = 10 ⁻¹⁵ , EW15	≥ 0.46 (host to module) ≥ 0.57 (module to host)
CEI-28G-VSR: Eye height at BER = 10 ⁻¹⁵ , EH	≥ 95 (host to module) ≥ 228 (module to host)
CEI-28G-VSR: Vertical Eye Closure (VECP)	≤ 5.5 dB (module to host)
CEI: Capable of Total Jitter at BER = 10 ⁻¹⁵	TJ: ≤ 0.28-0.54 UI

Table 5. Summary of typical electrical transmitter requirements.

3.3. Testing electrical transmitters

The broad range of electrical transmitter requirements given in Table 5 demonstrates the differences in the electrical signaling subsystems shown in Figure 2: serdes to serdes, serdes to

transceiver, transceiver to serdes across cables or a backplane with different connector configurations. Test requirements vary with application. Introduction of FEC, loosens constraints where it's mandatory, for 100 GbE, but not where it's optional, for CEI.




Figure 12. Examples of insertion loss response, $IL(f)$, for a few compliance test boards/ISI stressors.

Configure the transmitter being tested to **transmit the appropriate test pattern**. PRBS9 is usually sufficient. PRBS31 or RS-FEC scrambled idle will always suffice, but check the specification to see if they're overkill. Of course, all other system channels should be active with specified peak-to-peak voltages and rise/fall transition times so that crosstalk is included in the tests. Crosstalk aggressors should transmit different patterns and, to the extent possible for the reasons discussed above, the aggressors should be asynchronous.

Since eye diagrams at these data rates, even after transmission across just a few cm of PCB or cables, can be closed at the receiver, the transmitters must apply equalization in the form of pre- or de-emphasis, a type of FFE (feed forward equalization).

At least **three taps of pre-emphasis** are required. Three taps means that the voltage levels of a bit making a transition as well as the preceding and following bits are modified to compensate for the channel frequency response. The

tap values, C_{-1} , C_0 , C_1 , are derived from the channel pulse response. Think of the ISI introduced by the channel as the effect of folding the channel frequency response over the profile of each transmitted bit. The resulting shape of each bit can extend over many UI. Typical 25+ Gb/s specifications require optimization over at least 8 UI.

Connect the test signal to the **appropriate compliance test board or filter**. Test boards are necessary to assure that transmitted signals are compatible with both compliant channels and the equalization capabilities of compliant receivers. Each specification requires compliance boards with different loss and frequency response profiles to emulate their application and length of transmission on PCB. A few insertion loss response examples are shown in Figure 12. It's important to keep in mind that the frequency response requirements apply to the *combination* of both the compliance test board and whatever test fixture you use to connect the transmitter under test to the board and the board to your test equipment: don't forget the test fixture insertion loss!

Rather than use test boards that require extensive S-parameter characterization, Tektronix offers a convenient and flexible alternative. The LE320 filter is an adjustable FFE. Just as an FFE can be used to correct the transmission path response, it can also be used to apply that response. Thus, the LE320 filter can replicate most compliance test boards.

With the appropriate pattern and compliance board in place, **connect the test equipment**. Transmitter characteristics can be measured on DSA8300, PatternPro Error Detector, or BERTscope triggered with a clock recovered from the signal. The CR286A clock recovery unit serves as the reference receiver with a golden PLL. Check the specifications to make sure you apply the correct 3 dB bandwidth.

CEI-28G-VSR is one of the more challenging standards. It has separate requirements, including separate compliance test boards, for host-to-module and module-to-host tests. CEI-28G-VSR also specifies maximum allowed **common mode noise** which is measured on the DSA8300 by summing, rather than subtracting, the differential signals. With a free trigger, accumulate a vertical histogram over 1 UI; the common mode noise is given by the square root of the histogram rms squared minus the square of the instrument noise.

Optimize the transmitter pre-/de-emphasis taps. Some specifications also require that the test equipment apply a CTLE and perhaps even a DFE so that the interaction of transmitter pre-emphasis and receiver equalization is included in the test. The CTLE is typically a single zero, two-pole filter that peaks at the Nyquist rate, $f_{data}/2$.

With the compliance board in place and pre-emphasis optimized. Configure the CTLE gain, typically 1 to 3 dB, if required by the spec to produce the greatest eye height. If you're using the DSA8300, collect at least 12 million samples; if you're using a PatternPro Error Detector or BERTscope, acquire at least 4 million bits. The greater the statistical sample, the better.

Compliant transmitters must pass requirements on several quantities: **EOJ** (even-odd jitter), a type of **DCD** (duty-cycle distortion), is the difference between the average deviations of even-numbered and odd-numbered logic transitions; **UUGJ** (uncorrelated unbounded Gaussian jitter) is, for practical purposes, equivalent to RJ; **UBHPJ** (uncorrelated bounded high probability jitter)—the DJ components that are not correlated to the data pattern—is primarily periodic jitter and jitter from crosstalk; **EBUJ** (effective bounded uncorrelated jitter); and, **ETUJ** (effective total uncorrelated jitter).

Since the eye is closed by DDJ (data-dependent jitter, the combination of ISI and DCD) caused by the frequency response of the channel, and since compliant receivers correct that DDJ, rather than specify TJ (total jitter) and EH, the 100 GbE electrical transmitter specifications, 100GBASE-28CR and KR, require maximum ETUJ (total uncorrelated jitter) at $BER = 5 \times 10^{-5}$ which measures the eye opening independent of DDJ.

The “effective” E prefix denotes that they are derived by a specific implementation of the dual-Dirac model. **ERJ** (effective RJ) completes the set of dual-Dirac parameters so that $ETUJ(BER) = Q \times ERJ + EBUJ$, with $Q = 7.9$ for $BER = 5 \times 10^{-5}$ and $Q = 14.1$ for $BER = 10^{-12}$.

In CEI-28G-VSR, the transmitted signal amplitude specification is given in terms of **eye height and eye width**. As above, EH and EW are both defined with respect to BER: $EHN = EH(10^{-N})$ and $EWN = EW(10^{-N})$. EHN and EWN are easy to measure with a BER-eye diagram on the DSA8300 equipped with 80SJNB BER analysis software or a BER Contour on the PatternPro Error Detector or BERTscope.




Figure 13. Measurements of EH15 and EW15 with a BER contour.

EH15 is the vertical separation of the inner $BER=10^{-15}$ contours at the center of the eye, as shown in Figure 13. Since direct measurement of EH15 can take hours, the value can be extrapolated from measurements of a few billion bits either automatically or by hand in a compliant way: To get EH15, measure the random noise on the high and low voltage rails. The rms deviations of vertical histograms at the center of the eye diagram give RNO and RN1. Then measure EH6 on the PatternPro Error Detector or BERTscope to get EW15 from $EH15 = EH6 - 3.19 \times (RNO + RN1)$.

CEI-28G-VSR **VEC** (vertical eye closure) is the ratio of the average voltage swing to eye height:

$$VEC = 20 \log \frac{V_{avg\ swing}}{EH15}$$

On the PatternPro Error Detector or BERTscope, you can read the average voltage swing right off the BER-contour. On the DSA8300, you have to acquire an eye diagram with at least 12 million samples to get reliable average high and low logic levels. Their difference is the average voltage swing.

Compliant CEI transmitters meet the specified EH at $BER=10^{-15}$, TJ at $BER=10^{-15}$, and VEC requirements given in Table 6.

Typical Electrical Receiver Compliance Criteria													
Stress pattern	PRBS31 or RS-FEC encoded scrambled idle												
Receiver jitter and noise tolerance tests													
Swing voltage	600-1200 mV												
4th order, low pass, Bessel-Thompson filter, 3 dB bandwidth	33 GHz												
Compliance test board or FFE filter with frequency response, $IL(f)$, consistent with specified parameters in $IL(f) = a_0 + a_1\sqrt{f} + a_2f + a_3f^2 + a_4f^3$ (dB)	For example, Figure 12 and Figure 15												
SJ at frequency above roll off in Figure 7	0.05 UI												
RJ (at BER = 10^{-15})	0.15 UI												
UBHPJ	0.15 UI												
Add more RJ and sinusoidal interference so that	EH(10^{-15}) = 240 mV and TJ(10^{-15}) = 0.43 UI												
VEC	4.5-5.5 dB												
Three tests, one for ideal CTLE and +1 dB and -1 dB peaking													
Compliance criteria (depending on standard): BER for any SJ amplitude and frequency in Figure 7 or BER without FEC or BER prior to FEC or FLR (frame loss ratio) post RS-FEC	$\leq 10^{-15}$ or $\leq 10^{-12}$ $\leq 5 \times 10^{-5}$ $\leq 6.2 \times 10^{-10}$												
Interference tolerance tests													
Both of two separate compliance test boards or FFE filters with frequency response, $IL(f)$, with parameters specified	$IL(f) = a_1\sqrt{f} + a_2f + a_4f^2$ (dB) <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th></th> <th>Test 1</th> <th>Test 2</th> </tr> </thead> <tbody> <tr> <td>a_1</td> <td>1.7</td> <td>4.3</td> </tr> <tr> <td>a_2</td> <td>0.546</td> <td>0.571</td> </tr> <tr> <td>a_4</td> <td>0.01</td> <td>0.04</td> </tr> </tbody> </table>		Test 1	Test 2	a_1	1.7	4.3	a_2	0.546	0.571	a_4	0.01	0.04
	Test 1	Test 2											
a_1	1.7	4.3											
a_2	0.546	0.571											
a_4	0.01	0.04											
SJ	0.1 UI at ≥ 100 MHz												
RJ (UUGJ)	0.01 UI (rms)												
DCD Even-odd jitter (DCD)	0.035 UI												
COM	≥ 3 dB												
Pass if RS-FEC symbol error ratio	$\leq 10^{-4}$												

Table 6. Summary of serdes and transceiver electrical tolerance specifications.

3.4. Testing electrical receivers

The essential idea of receiver testing is to subject the receiver to the worst case signal. If it operates at an acceptable BER, then it should be able to perform with any signal it's likely to see in the field. The introduction of FEC, crosstalk from parallel lanes, and the desire to make designs flexible complicate receiver testing; though the ultimate goal is to assure interoperability at BER $\leq 10^{-12}$ for 100 GbE and BER $\leq 10^{-15}$ for CEI.

Receiver testing includes noise and jitter tolerance tests, interference tolerance tests, or both. Each specification requires different levels and types of stress. In this section, we include a complete cross-section of the stresses you can expect to encounter in 100G technology with typical levels quoted in Table 6. Be sure to check the specification to which you are testing to guarantee that your test is compliant.

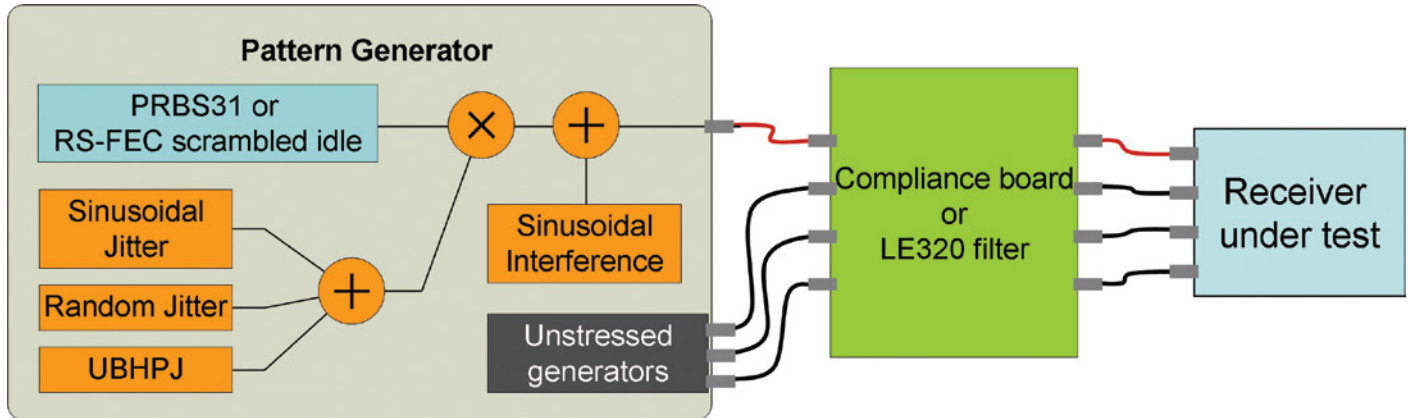


Figure 14. Electrical stressed receiver test setup.

To assemble the test, first set the appropriate **low pass filter** in the pattern generator. Then connect the compliance test board as shown in Figure 14, but with a PatternPro Error Detector, a BERTscope, or a DSA8300 in place of the receiver under test.

The combination of the **compliance test board** and the cabling of your test fixture must satisfy specific frequency response criteria. The specifications provide instructions for how to fit measured S-parameters to a parameterization like

$$IL(f) = a_0 + a_1\sqrt{f} + a_2f + a_3f^2 + a_4f^3 \text{ (dB)}$$

To achieve the correct frequency response, first measure the S-parameters of the combined test board and test fixture. Then calculate the fit parameters and compare to those required by the specification.

A convenient alternative to measuring and fitting S-parameters, is to simply insert a modular FFE filter, Tektronix's LE320, between the test fixture and the device under test. Then set the FFE taps to apply the specified frequency response.




Figure 15. Insertion loss response curves, $IL(f)$, for the two interference tolerance tests required of 100 GbE 100GBASE-CR4 cables.

To apply different levels of ISI, some standards require testing with more than one compliance test board. CEI-28G-VSR has separate compliance test boards for the host-to-module and module-to-host directions. 100GBASE-CR4 requires two separate interference tolerance tests, each with a specific compliance board; the $IL(f)$ responses for the two are shown in Figure 15.

Configure the stressed signal, as shown in Figure 14. Set the signal amplitude to the specified level, usually in the range from 600-1200 mV; the latter extreme tests the receiver's input amplitude tolerance. Transmit the stress test pattern, either a PRBS31 or RS-FEC encoded scrambled idle pattern. Both are long patterns capable of producing every reasonably likely ISI impairment.

Optimize the transmitter pre/de-emphasis taps.

Apply crosstalk by engaging pattern generators on the other channels. If it's not possible for each aggressor to transmit a unique pattern, it's imperative that you introduce sufficient delay between them so that the patterns aren't synchronized or you risk failing a compliant receiver.

If you're testing to a standard that prescribes EH and EW at the receiver input, like CEI-28G-VSR, then **apply 0.05 UI of SJ**. At this point, you're still setting up the stress signal, so just apply that 0.05 UI of SJ at a frequency above the roll off in Figure 7. You'll stress the receiver's clock recovery circuit by varying the SJ frequency and amplitude later. **Apply 0.15 UI of UBHPJ** and 0.15 UI **UUGJ** or, equivalently, RJ. The pattern generator applies UBHPJ (also known as BUJ) by mixing a low rate PRBS pattern with the signal. Now, apply both **SI** (sinusoidal interference) and more RJ until you get the prescribed EH and TJ values. The resulting VEC should be in the range 4.5-5.5 dB. You can see the impact of the stress on the signal in Figure 16.

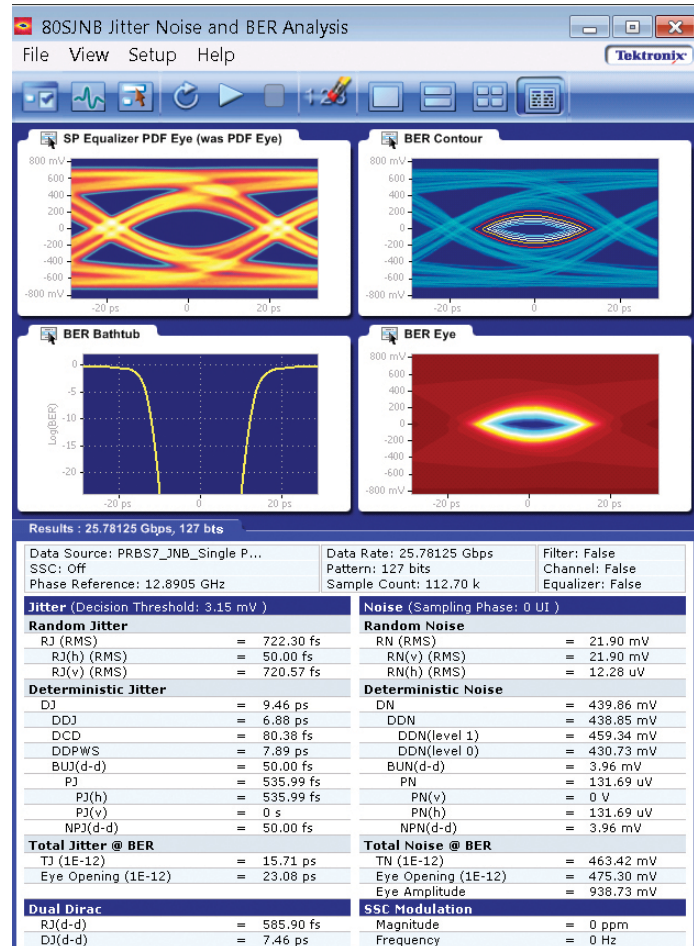


Figure 16. A stressed signal.

The tolerance test for systems that specify EH and TJ (or EW) prior to receiver equalization is now ready. Since the actual testing is the same for these systems as the rest, we'll get back to it after we finish setting up the others.

If you're testing one of the other CEI standards, CEI-28G-SR, CEI-28G-MR, or CEI-25G-LR or an electrical 100 GbE system, 100GBASE-CR4 with cables or a 100GBASE-KR4 with a backplane, the eye is quite likely closed at the receiver so EH and TJ (or EW) don't make sense until after receiver equalization. In these cases, **apply RJ, BUHPJ, and EOJ** (even-odd jitter) at the specified stress levels.

If **COM** (channel operating margin) is specified, the crosstalk aggressors should conform to the standard, $COM < 3$ dB, but the stress signal will have $COM \geq 3$ dB.

The parameter, **RSS_DFE4**, in Table 6, is the root sum square of the first four cursors of the impulse response. In other words, RSS_DFE4 is a figure of merit for how much the receiver DFE must impact the signal to open the eye.

To perform the test, do whatever you can to assure that the receiver sees the signal you've configured. It's best to connect the receiver being tested to the compliance board with the same fixture used to configure the stressed signal.

If your test receiver uses a training sequence to set its equalization parameters, first run the training pattern and optimize the receiver equalizer.

Since RS-FEC provides the ability to correct and/or count errors—and count more errors than it can correct—most systems we're addressing should be able to report their own FLR (frame-loss ratio) after applying FEC, BER prior to FEC, and/or RS-FEC symbol error ratio. If your system doesn't provide some form of error counting, then connect the receiver output to a Pattern Pro Error Detector or BERTscope. If the receiver doesn't provide a clock output, you're supposed to use a clock recovery unit to time the error detector. If you don't have a clock recovery unit, the pattern generator's data rate clock should suffice unless you're operating over a long distance link.

Apply the stressed signal to the receiver first with low amplitude SJ applied above the roll off frequency, Figure 7. If the receiver operates with its equalization scheme enabled and optimized at or better than the specified BER, FLR, or symbol error ratio, continue testing across the SJ frequency-amplitude template in Figure 7 to assure that the receiver can track low-frequency jitter with all other stresses applied. It's automatic with the jitter tolerance measurement functions of the PatternPro Error Detector or BERTscope.

The receiver is compliant if it operates at or better than specified across the SJ frequency range.

4. Diagnostic tests

The difference between compliance and diagnostic tests is complexity. Compliance tests include too many elements to resolve the problem when something goes wrong. To determine the components of a system that cause problems, diagnostic tests should be strategically planned to probe specific weaknesses. They should build in complexity, test upon test, to find problems and determine margins.

4.1. What to do if the transmitter fails

If the transmitter fails, simplify the test conditions by removing any test compliance boards and analyzing the transmitter output with as direct a connection as possible. Perform jitter and noise analysis. Analyze the breakdown as you apply more complex patterns. Increase ISI gradually with test boards that have increasing insertion loss or the equivalent by adjusting the LE320 filter taps, apply pre-emphasis, and turn on crosstalk aggressors. For each set of conditions, analyze eye diagrams, BER-eye, BER-contours and the jitter and noise breakdown, Figure 17.

The jitter map automatically distinguishes different types of jitter to help isolate problems:

- Non-periodic BUJ → crosstalk isn't being sufficiently shielded.
- DCD → transmitter distortion.
- ISI → trouble with the output path or package.
- High levels of RJ → problems with transmitter clocking.
- Sinusoidal and periodic noise and jitter → electromagnetic interference from a nearby component, perhaps a switching power supply. Identify the source of interference by studying the jitter frequency spectrum. The clock recovery unit, CR286A, measures the real time jitter spectrum. Do the frequencies of any spectral peaks correspond to harmonics of other components?




Figure 17. Example of jitter breakdown.

4.2. What to do if the receiver fails

Investigate the receiver's response to each stress. Start with a clean pattern and add complexity and use the features of the pattern generator.

- Apply test patterns with low mark density to check for baseline wander.
- Challenge clock recovery circuits with patterns that have long CID (consecutive identical) bit strings and low transition density. The more structure a pattern has, the more ISI it generates when applied to a compliance board or filter.
- Analyze the locations of detected errors to correlate the problem to specific pattern sequences or other deterministic effects using the BERTScope Error Analysis functions.
- Sweep SJ across the receiver's clock recovery frequency response at different amplitudes. Determine where the clock recovery circuit's ability to track jitter breaks down.
- Challenge the receiver's equalizer. Introduce test boards with increasing insertion loss or the equivalent by adjusting the LE320 filter taps. By combining pattern complexity and trace length, you can generate a wide variety of ISI levels and find an equalizer's margin.
- Probe the receiver's ability to tolerate jitter and noise, that is, its setup and hold, by ramping up crosstalk and RJ.
- Check the voltage sensitivity by introducing sinusoidal interference.

Find the sensitive aspects of the receiver and then apply different stress combinations. There may be stress combinations that are particularly challenging as well as peculiar combinations where the receiver is quite robust.

Summary

Signal transmission at 100 Gb/s is not simple. You need high performance test equipment to evaluate the performance of every component and the system as a whole. Between the PPG PatternPro Pattern Generators and PED Error Detectors, the DSA8300 equivalent-time sampling oscilloscope with 80SJNB software, the BSA286 BERTscope, CR286A clock recovery units, and LE320 FFE filter, Tektronix provides a complete set of high performance instruments for both optical and electrical transmitter and receiver compliance and diagnostic testing.

Contact Tektronix:

- ASEAN / Australia** (65) 6356 3900
- Austria*** 00800 2255 4835
- Balkans, Israel, South Africa and other ISE Countries** +41 52 675 3777
- Belgium*** 00800 2255 4835
- Brazil** +55 (11) 3759 7627
- Canada** 1 (800) 833-9200
- Central East Europe and the Baltics** +41 52 675 3777
- Central Europe & Greece** +41 52 675 3777
- Denmark** +45 80 88 1401
- Finland** +41 52 675 3777
- France*** 00800 2255 4835
- Germany*** 00800 2255 4835
- Hong Kong** 400-820-5835
- Ireland*** 00800 2255 4835
- India** +91-80-30792600
- Italy*** 00800 2255 4835
- Japan** 0120-441-046
- Luxembourg** +41 52 675 3777
- Macau** 400-820-5835
- Mongolia** 400-820-5835
- Mexico, Central/South America & Caribbean** 52 (55) 56 04 50 90
- Middle East, Asia and North Africa** +41 52 675 3777
- The Netherlands*** 00800 2255 4835
- Norway** 800 16098
- People's Republic of China** 400-820-5835
- Poland** +41 52 675 3777
- Portugal** 80 08 12370
- Puerto Rico** 1 (800) 833-9200
- Republic of Korea** +822-6917-5000
- Russia** +7 495 664 75 64
- Singapore** +65 6356-3900
- South Africa** +27 11 206 8360
- Spain*** 00800 2255 4835
- Sweden*** 00800 2255 4835
- Switzerland*** 00800 2255 4835
- Taiwan** 886-2-2656-6688
- United Kingdom*** 00800 2255 4835
- USA** 1 (800) 833-9200

* If the European phone number above is not accessible,
please call +41 52 675 3777

Contact List Updated June 2013

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2015, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

