

InfiniBand Compliance Testing with Real-time Oscilloscopes



▶ Introduction

InfiniBand promises to bring channel-based I/O reliability and performance to the world of distributed computing, but with it comes ultra-high signal rates that challenge conventional test and measurement solutions. High-performance, real-time oscilloscopes can answer that challenge by combining specific analysis features and practical versatility with the world's fastest real-time acquisition technology.

As processor speeds increase and the Internet drives the demand for constantly available data, Internet data centers – the "mission controls" of information processing – face a rise in complexity. The biggest challenge to improved data center performance is the I/O subsystem.

InfiniBand™ architecture resolves this challenge by providing a mechanism – a unified fabric – to share I/O interconnects among many servers. This architecture creates a more efficient way to connect storage and communications networks and server clusters together, while delivering an I/O infrastructure that offers the performance, reliability, scalability and flexibility required by data centers. It provides scalable performance through multi-link connections, specifying point-to-point connections with a 2.5 Gb/s wire speed and offers three levels of link performance – 1X (2.5 Gb/s), 4X (10 Gb/s) and 12X (30 Gb/s).

InfiniBand Testing Basics

InfiniBand compliance and interoperability testing is rapidly being formalized into a rigid framework of pass/fail testing procedures.

These procedures cover everything from link layer and data management issues to low-level physical layer tests. In the InfiniBand architecture specification, Chapter 6 (High-Speed Electrical Signaling) defines the high-speed electrical requirements, while Chapter 7 (Copper Cable Specifications) covers cable test requirements and Chapter 8 (Fiber Attachment) specifies optical measurements. InfiniBand fiber requirements are derived from SONET OC-48/SDH STM16 standards and InfiniBand electrical specifications are based on Fibre-Channel (ANSI NCITS/1235D). Building upwards from these two solid foundations gives InfiniBand a head start on developing a cohesive standard and encourages test and measurement manufacturers, such as Tektronix, to leverage their experience with these foundation standards.

It is critical in compliance and interoperability testing to eliminate any opportunities for multiple interpretations of results or confusing test procedures. This is especially difficult when dealing with very high-speed signal rates because one must be intimately familiar with working beyond the classical digital abstraction layer and be prepared to apply microwave analysis techniques where necessary.

InfiniBand Compliance Testing

► Application Note

If there was one major lesson to be learned from conducting similar testing in the Fibre-Channel and SONET arenas, it is that compliance test methodology must be clear and concise enough to avoid requiring dissertations on equipment used and interconnect techniques employed.

A Real-time Solution

Previously, engineers in multi-gigabit serial communications were limited to equipment designed for high-end telecommunications applications, such as equivalent time based communications signal analyzers and bit error rate testers. These instruments lack the versatility, utility and ease-of-use needed by these engineers.

Today's high-performance, real-time oscilloscopes, such as the TDS7000, CSA7000 and TDS6000 Series, can now perform physical layer serial bit sequence testing of InfiniBand traffic, enabling engineers to capture this traffic with enough resolution to visualize and conduct meaningful physical layer measurements on the signal. These instruments have the capability to make efficient InfiniBand interoperability measurements. They can support a broad spectrum of high-speed physical layer measurements that range from communication mask testing to comprehensive jitter analysis. With their high bandwidth and ease-of-use, they help designers and test engineers to achieve stable, repeatable test results, making them ideal design, troubleshooting, and compliance validation tools.

System Bandwidth Considerations

Physical layer analysis of InfiniBand traffic requires test and measurement instrumentation with at least 4.0+ GHz analog bandwidth. (This requirement follows the Fibre Channel committee's recommendation of a ratio of at least 1.8 between the bit rate and an oscilloscope's bandwidth.) Few test and measurement solutions can meet this requirement. To operate at these aggressive signal speeds and to conduct signal integrity measurements on multi-gigabit serial communications rates poses a formidable challenge for any instrument.

The rule of thumb recommended by the Fibre-Channel T11.2 working group for acquisition system bandwidth is 1.8X the bit-rate. This bandwidth recommendation is intended to capture the signal's entire spectral content up to about midway between the third and fourth harmonic. Therefore InfiniBand, with its 2.5 GHz bit-rate (1.25 GHz first harmonic), requires a 4.5 GHz analog bandwidth from probe tip to acquisition. For most measurements, however, a 1.5X rather than the 1.8X factor is adequate since the .3 difference is largely for error tolerance.

Sampling oscilloscopes have traditionally been strong candidates for many of these measurements since they typically have more than 40 GHz of analog bandwidth. However, the usefulness of sampling oscilloscopes is seriously limited by their need for a repetitive, deterministic stimulus to render even the most basic waveforms. The requirement for a Frame or Synchronous trigger when using this class of instrumentation makes them particularly effective for device characterization, but difficult to use as debug tools.

High-performance, real-time oscilloscopes, like the TDS6604 and TDS7404, are a more flexible choice for general high-speed applications. The TDS6604 offers up to 6 GHz bandwidth (fast enough even for 3.2 Gb/s XAUI measurements) and sample rates on two channels of up to 20 GS/s. This performance allows users to sample both differential lanes, with approximately 8 sample points per unit interval (bit time) and gives reasonable assurance of finding two or three sample points per rising and falling edge, as well as sample points over the intervening high or low bit interval.

Signal Trigger Requirements

8B/10B encoding is a technique whereby 8 bits of logical data is expanded to 10 bits of physically transmitted data. These additional 2 bits are strategically placed in the outbound physical data stream to redistribute the signal's peak spectral content around the first harmonic. The first harmonic is also the bit rate of the transmitter and receiver clock. The receiver on the other end of the physical cable utilizes this encoded clock information to keep its phase locked loop time-base synchronized with the incoming data stream.

To properly trigger an oscilloscope on this form of signal stream requires the use of a Clock Data Recovery unit (CDR) to separate the clock and data information. In this sense, the TDS6604 or CSA/TDS7404 oscilloscope acts as a sophisticated InfiniBand Receiver and can lock onto data with a synthesized bit rate clock that tracks with the incoming signal clock. An instrument with a CDR trigger capability can also expose pattern dependent anomalies and dropouts that are easily missed with edge detection, and provides a stable integrated trigger source with good low frequency jitter rejection.

Probing and Interconnects

Another challenge is to get the signal of interest to the testing instrument without loading it down or distorting its data. Current InfiniBand compliance test methodology requires breaking the InfiniBand link and relying on the transmitter's broadcast of a beacon

signal as the basis for many of the compliance measurements. Using an InfiniBand-to-SMA breakout (as shown in Figure 1) channels the 50-Ohm signal paths directly into the 50-Ohm termination of the oscilloscope. This is ideal both from a termination and a practical standpoint. For example, the break out board allows an easy transition from channel to channel to permit quick testing of all 24 lanes of a 12X cable.



► **Figure 1.** InfiniBand-to-SMA breakouts allow you to connect the 50-Ohm signal paths directly into the 50-Ohm inputs of the oscilloscope.

Directing the InfiniBand signal paths into the TDS6604 allows its PLL-based CDR option to emulate the actual InfiniBand receiver, both in its tracking characteristics and in its electrical termination. The CDR supports user-programmable tracking and configures itself, by default, as a "Golden PLL" as defined by Fibre Channel T11.2. (A Golden PLL has loop bandwidth tracking characteristics that allow it to reject low frequency aberrations that frequently occur below 1.5 MHz, or bit rate/1667.)

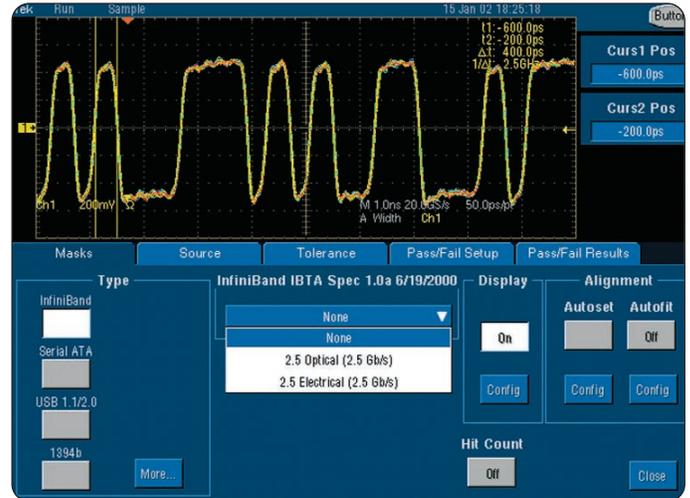
Typical InfiniBand Measurements

Having reviewed the basics of the InfiniBand technology and the tools needed for its compliance testing, it is time to consider some typical InfiniBand high-speed electrical measurements.

AC Parametric Measurements

AC Parametric measurements can be made with any real-time oscilloscope of sufficient bandwidth. By using a host channel adapter in a beaconing (Polling) state, the measurements are well within the capabilities of most oscilloscope measurement packages.

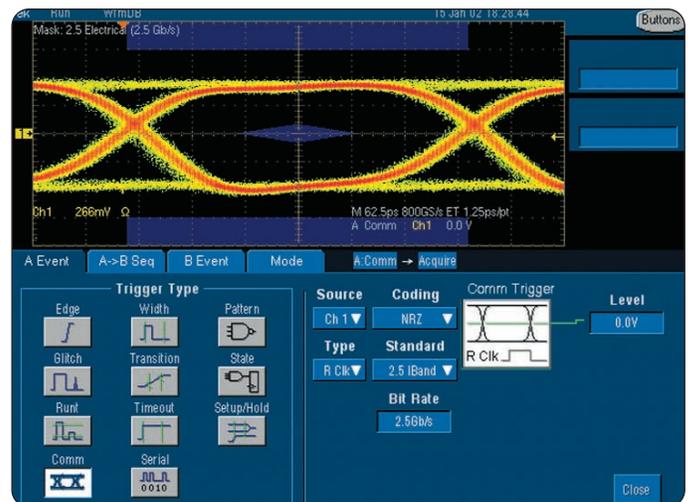
To examine a TS1 packet, for example, use a pulse width trigger off of the Comma character (5 consecutive pulse unit intervals), and then apply a trigger delay to index into areas of interest in the information packet (see Figure 2).



► **Figure 2.** A TS1 packet at 400 ps using a pulse width trigger.

Eye Diagram Measurements

Eye diagram measurements are easy to make with TDS6000, TDS7000 and CSA7000 Series digital real-time oscilloscopes. To examine the same TS1 packet as an eye diagram, simply use the oscilloscope's mask test feature to select the specification standard you want to use. Then press auto-set. The oscilloscope automatically switches to an eye diagram rendering of the same signal.



► **Figure 3.** An eye diagram rendering of the same TS1 packet.

InfiniBand Compliance Testing

► Application Note

When rendering an eye diagram, the oscilloscope automatically switches the triggering system to CDR mode and specifies the proper bit rate to lock on and track the signal. In our example (see Figure 3), we can see an uncharacteristically clean reference signal that illustrates the precision of the CDR system. Note we have a RMS trigger jitter of about 6 ps with a Pk-Pk trigger jitter of around 42 ps.

Users can adjust the CDR loop tracking characteristics to custom rates or switch to one of the other triggering modes supported by the TDS6604. For example, if one needs to selectively visualize other portions of the eye diagram, set the pulse width trigger to react to one of the five consecutive pulse fiducial found in the K28.5's. Call on the trigger delay index to display different parts of the pattern sequence. Another way to use the TDS6604 is to assess the equalization function by comparing the eye diagrams of short and wide pulses.

Eye Diagram Testing with the TDS6604

Eye Diagram tests are a major component of the compliance and interoperability measurements listed in the InfiniBand specifications. Specifically, the tests detailed in Chapter 6, section 6.5.1 for receiver input compliance tests, require eye diagram measurements. The TDS6604 real-time oscilloscope is an excellent choice for these measurements. To perform an Eye Diagram test on an electrical signal, follow these general steps.

(1) Connect the signal from the far end of the transmitter, exactly as it would enter the receiver, into the TDS6604 by using a breakout board. Start by looking at a single line of the differential pair, or use a differential probe such as the P7330. An InfiniBand waveform will appear on that channel. Select Mask: InfiniBand from the mask menu. (Note: This procedure assumes that there is an Option SM on the TDS6604.)

(2) Select 2.5 Gb/s Electrical.

(3) Press Auto Set.

(4) The eye diagram will appear on screen as shown in Figure 3.

(5) If one needs further analysis or measurement, several of the standard trigger modes are effective on signals at this rate (e.g., pulse width, timeout, and glitch). One can leverage these trigger modes to narrow the instruments trigger onto specific areas of interest.

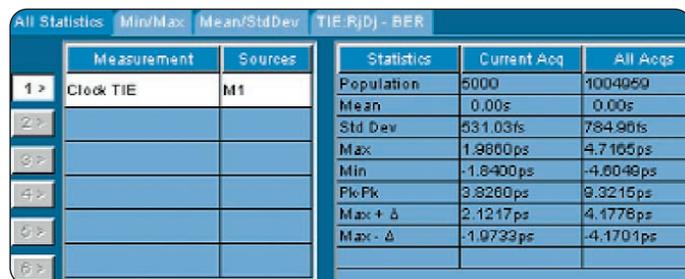
(6) It is also easy to use the native instrument mask testing capabilities of the TDS6604 to run a known number of waveforms through the mask testing system. The system can then to stop and report a PASS or FAIL condition after it processes the prescribed waveform population.

Jitter Analysis

At high data speeds, jitter can be a major source of error.

The InfiniBand Electrical Signaling specification calls out stringent jitter timing budgets for InfiniBand devices (Sections 6.4.3 and 6.5.1 of the InfiniBand Architecture Specification). To be sure of compliance, the tester must break the jitter down into its constituent components, namely Deterministic Jitter (Dj) and Total Jitter (Tj). Most InfiniBand jitter specifications require TIE (Time Interval Error) jitter measurements. Their principal focus is to determine how much the signal deviates from the precise theoretical time transitions that the encoded clock would dictate in the signal stream.

Identifying and measuring the jitter components allows you to debug the circuits and ensure that the jitter stays within set limits. Previously, this required using a slow, specialized (and expensive) Bit Error Rate Tester (BERT). Now, TDSJIT3 jitter analysis, coupled with a Tektronix high-performance, real-time oscilloscope, offers the superior jitter solution. This solution uses clock recovery algorithms that rely on a constant clock rate or on a Golden PLL to recover ideal timing transition information from the signal stream. TDS7000 or CSA7000 Series oscilloscopes, equipped with TDSJIT3 software, calculate the time interval error or jitter over long record lengths and separates the results into Random (Rj) and Deterministic (Dj) elements. TDSJIT3 software in such a system can determine Bit Error Ratio results much faster than even a dedicated BERT can. If higher-frequency jitter components are of interest, the TDS6604 allows the analysis of jitter results down to 700 fs with 6 GHz system bandwidth (see Figure 4).



All Statistics		Min/Max	Mean/StdDev	TIE/Rj/Dj - BER	
	Measurement	Sources	Statistics	Current Acq	All Acqs
1 >	Clock TIE	M1	Population	5000	1004959
2 >			Mean	0.00s	0.00s
3 >			Std Dev	531.03fs	784.96fs
4 >			Max	1.9860ps	4.7165ps
5 >			Min	-1.8400ps	-4.6048ps
6 >			Pk-Pk	3.8260ps	9.3215ps
			Max + Δ	2.1217ps	4.1778ps
			Max - Δ	-1.0733ps	-4.1701ps

► **Figure 4.** Measurement of a very low Jitter Noise Floor (JNF). Note the Peak-to-Peak of 3.8 ps and the Standard Deviation of 531 fs RMS.

The measurement shown used a high precision BERT (less than 1 ps RMS jitter) for a signal source. The RMS jitter observed by the oscilloscope is 531 fs over a 5K-edge population. For populations over a million, the RMS jitter remains below 800 fs. The peak-peak jitter does not exceed 10 ps.

Bit Error Rate (BER) Measurements

BER measurements are necessary for a variety of serial data standards, including Fibre Channel and InfiniBand. The ability to perform BER measurements rapidly is important for component characterization as well as manufacturing test. But regardless of speed, the necessity to make accurate measurements remains paramount. To this end, the TDSJIT3 software package has been tested and correlates to within 2.1% of standard BERT testers at a 10E-12 BER. In measured results, where no estimation is used, TDSJIT3 correlates to BERT testers within 1.1%. TDSJIT3 performs Rj/Dj separation and BER measurements quickly with a very high degree of accuracy.

Resolving Jitter Components

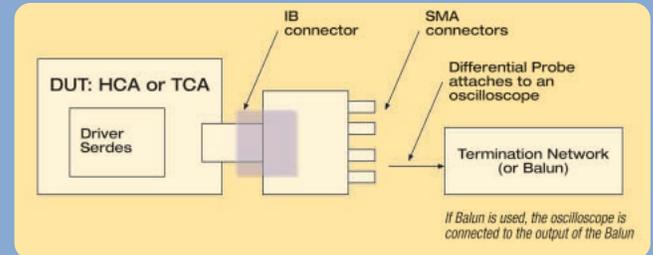
Jitter first became necessary as a signal quality specification when data rates reached the 200-500 Mb/s range. Now, with InfiniBand signals at 2.5 Gb/s, jitter testing is crucial. Using TDSJIT3 software with TDS6000, TDS7000 or CSA7000 Series oscilloscopes makes it simple to measure Total Jitter (Tj) and analyze its contributing components. This package characterizes Random Jitter (Rj) by its Gaussian distribution and divides Deterministic Jitter (Dj) into Periodic Jitter (Pj), Duty Cycle Distortion (DCD), Data Dependent Jitter (DDj) or Inter Symbol Interference (ISI). Once the jitter components are isolated and quantified, the package delivers an accurate estimate of the Bit Error Rate.

Testing Driver Jitter with a Real-Time Oscilloscope

Because jitter at high data speeds can be a major problem, the InfiniBand Electrical Signaling specification sets stringent jitter limits (Sections 6.4.3 and 6.5.1 of the InfiniBand Architecture Specification). This test example shows how to check driver jitter using a Tektronix high-performance, real-time oscilloscope and a differential probe. The test is defined for a single lane and must be repeated for each lane of the device under test (DUT).

To perform this test efficiently, employ the TDSJIT3 jitter analysis software package with a TDS7404 or CSA7404 real-time oscilloscope. To ensure maximum sample rate and an accurate jitter measurement, be sure to convert the signal from differential to single-ended by using a P7330 differential probe to connect across a termination network, or by using a Balun to convert the signal (see Figure 5). The example assumes the use of a TS1 (pattern length 320) as the test pattern. However, one can extend

the procedure to cover compliance jitter patterns such as CJTPAT (pattern length 2640). Regardless, the test pattern must be repetitive and have 100 or more repeats.



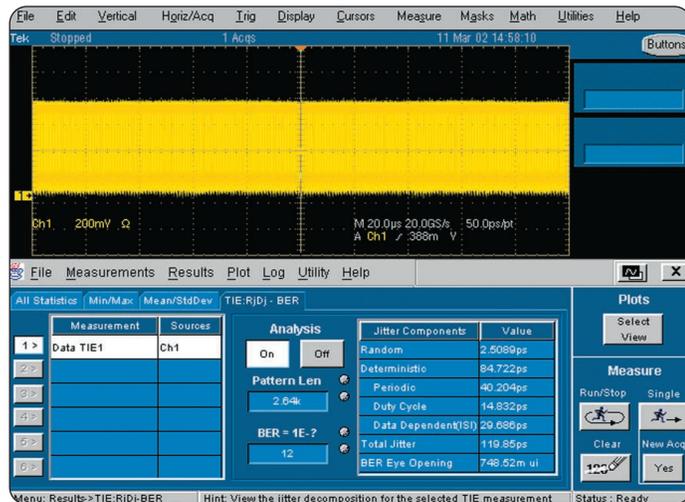
► **Figure 5.** Drawing of a typical Driver Jitter Test setup.

Initialize the test by sending the TS1 Bit Pattern. Then follow these general steps.

- (1) Probing Setup: Attach a PC1X_sma_P1 or similar test fixture to output of DUT. Then Attach the Differential Probe tip across terminated IBtxOP, IBtxON differential signal. Finally, attach the Differential Probe to Ch1 of the Oscilloscope.
 - (2) Oscilloscope Setup: Press AUTOSET. Then set the oscilloscope's horizontal timebase to 2 μ s per division. To ensure you have adequate timing resolution, set the sample rate to 20 GS/s or 50 ps per sample point. This will capture 20 μ s of real-time data (50,000 contiguous bit times, 156 pattern repeats). Remember that this setting will depend on the pattern length selected. Total required time captured must equal or exceed 100 times the pattern length multiplied by the time interval. For higher accuracy, the horizontal timing RESOLUTION can be increased to 10 picoseconds per point (ps/pt) using sin(x)/x interpolation mode. Select Measurement>Amplitude>Amplitude to verify that differential voltage level is within driver limits of 1.0 to 1.6 Vp-p.
 - (3) Measure Jitter using TDSJIT3 software: Select File > Run Application > Jitter Analysis 3 to run the test software. Select Data > TIE measurement using Data Source = Ch1; If the signal contains significant low frequency jitter or wander, then use PLL TIE measurement. Configuring Loop BW= IB2500: 2.5 sets the PLL loop bandwidth to (2.5G/1667)Hz, or 1.5 MHz. Perform Autoselect on Ref Levels. Go to Results and press the Run button to capture a Single acquisition.
- Choose the TIE:RjDj – BER tab in TDSJIT3 software; input a pattern length of 320 bits; and select Analysis = On. Deterministic jitter (Dj) and Total jitter (Tj) are listed as line items in the results display.

InfiniBand Compliance Testing

► Application Note



► **Figure 6.** The Compliance Jitter Test Pattern is the preferred signal for jitter analysis. This screen shot shows the results of the jitter decomposition discussed in the previous sections.

The Compliance Jitter Test Pattern

The Compliance Jitter Test Pattern is the preferred signal for jitter analysis. It covers a wide frequency and is able to stress receiving circuits to the extreme limits of their tracking capability (see Figure 6).

Instrumentation Requirements for InfiniBand Compliance Testing

Many different measurements are required by the InfiniBand specifications to test compliance. The sheer number of tests, as well as the high speed of the signals, makes testing a great challenge for engineers and designers. They need capable, versatile measurement instruments to help make their high-speed testing the precise, repeatable science it needs to be.

High performance, real-time oscilloscopes, such as the TDS6000, TDS7000 and CSA7000 Series, and digital sampling oscilloscopes, like the CSA8000B, can be paired with Tektronix high-speed probes, signal sources and test software to fill that need. Their superior measurement fidelity, powerful analysis capabilities, and flexible connectivity make them excellent choices for the measurement of high-speed designs.

TDS6604 Digital Storage Oscilloscope

► The TDS6604 DSO is a superior, high-performance solution for verification, debug and characterization of electrical signals for InfiniBand implementation. The oscilloscope features 6 GHz bandwidth and 20 GS/s sample rate that exceeds specifications required for InfiniBand compliance testing. The TDS6604, equipped with TDSJIT3 jitter analysis software, offers jitter measurements down to 0.7 ps RMS typical as well as a battery of optional compliance mask tests that quickly adapt it to the requirements of different stan-

dards. This instrument includes a complete parametric measurement system for signal characterization and allows user-defined math expressions to be performed on waveform data for customized on-screen results. Its mask testing with clock recovery (Option SM) enhances testing on serial data streams from 1.5 Mbaud to 2.5 Gbaud. The TDS6604's classic analog-style controls, touch sensitive display, and graphical menus make it exceptionally easy to learn and use. It offers an open Windows™ environment and built-in networking to encourage integration into sophisticated test and analysis applications.



CSA7404 Communications Signal Analyzer

► The CSA7404 communications signal analyzer, with its real-time digital phosphor oscilloscope architecture, meets the needs of designers who are challenged to implement emerging electrical and optical communications standards, such as InfiniBand, at unparalleled data rates up to 2.5 Gb/s. This instrument combines industry-leading speed with communications-focused capabilities to enable engineers to more efficiently design and debug optical and electrical communications equipment.

The CSA7404 features a real-time data capture rate of 20 Gb/s, a long record length up to 32 MB, patented third-generation DPX™ waveform capture technology, an advanced triggering system, and operates at up to 2.5 Gb/s. It can be used with user-controllable, full-featured, PLL-based CDRs on either the electrical or the optical inputs to provide stable trigger sources on clock-encoded data streams. This real-time communication signal analyzer delivers the flexibility and convenience of software-based optical receiver reference (ORR) filtering that supports a full range of telecommunications, data communications and computer serial bus standards in one instrument. It offers three operating modes: as a 4 electrical channel DPO; as a 3 electrical channel and 1 optical channel DPO; or as a one-channel optical reference receiver (ORR).



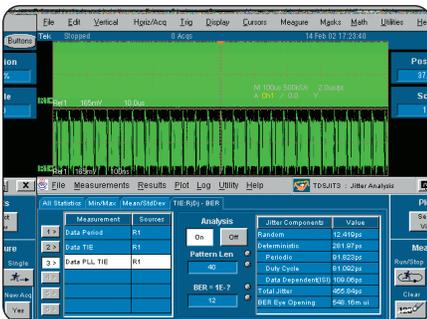
CSA8000B Communications Signal Analyzer

- Designed for high-performance communications applications, the CSA8000B communications signal analyzer, with its digital sampling oscilloscope architecture, is the best choice for characterizing InfiniBand optical signals. This instrument offers excellent bandwidth, time base, vertical accuracy, and jitter performance. It features bandwidth from DC to 50 GHz, automatic communication measurements, and automatic mask testing, and addresses data rates from OC-3/STM1 (1 Mb/s) to OC-768/STM256 (40 Gb/s). The CSA8000B provides built-in optical clock and data recovery, FrameScan™ to identify individual bits that fail mask testing, and TDR capabilities for device impedance characterization. It provides optional expansion up to 8 electrical channels. The CSA8000B's intuitive interface and color-graded display make it very easy to learn and use.



TDSJIT3 Jitter Analysis Software

- The TDS6000, TDS7000 and CSA7000 Series oscilloscopes, equipped with TDSJIT3 jitter analysis software, provides the highest accuracy jitter measurements available. With its comprehensive jitter analysis algorithms, it simplifies the discovery of jitter and its related sources in high-speed digital and communication systems. This software package enables engineers to characterize random and deterministic jitter (Rj/Dj) and predict bit error ratio (BER) to quickly determine overall system quality and rapidly deliver more robust designs to market.



P7260 Active Probe

- The P7260 single-ended active probe delivers an unprecedented 6 GHz bandwidth and < 75 ps rise time capability to enable engineers to view the fastest signals for their next-generation digital designs. Its high performance, low circuit loading and low noise make it an excellent choice for InfiniBand testing.



AWG710 Arbitrary Waveform Generator

- The AWG710 arbitrary waveform generator combines world-class signal fidelity with ultra high-speed mixed signal simulation. Its 4.0 GS/s sample rate simulates real-world signals up to 2.0 GHz. Its powerful sequencing capability, graphical user interface and flexible waveform editor make it a valuable companion in InfiniBand compliance testing.



Conclusion

InfiniBand™ architecture brings greater reliability and performance to distributed computing by creating a more efficient way to connect storage and communications networks and server clusters together, but with it comes ultra-high signal rates and a suite of high-speed compliance tests that challenge conventional test and measurement solutions.

High-performance, real-time oscilloscopes can answer that challenge by combining specific analysis features and practical versatility with the world's fastest real-time acquisition technology.

Contact Tektronix:

ASEAN Countries & Pakistan (65) 6356 3900

Australia & New Zealand (65) 6356 3900

Austria +43 2236 8092 262

Belgium +32 (2) 715 89 70

Brazil & South America 55 (11) 3741-8360

Canada 1 (800) 661-5625

Central Europe & Greece +43 2236 8092 301

Denmark +45 44 850 700

Finland +358 (9) 4783 400

France & North Africa +33 (0) 1 69 86 80 34

Germany +49 (221) 94 77 400

Hong Kong (852) 2585-6688

India (91) 80-2275577

Italy +39 (02) 25086 1

Japan 81 (3) 3448-3111

Mexico, Central America & Caribbean 52 (55) 56666-333

The Netherlands +31 (0) 23 569 5555

Norway +47 22 07 07 00

People's Republic of China 86 (10) 6235 1230

Poland +48 (0) 22 521 53 40

Republic of Korea 82 (2) 528-5299

Russia, CIS & The Baltics +358 (9) 4783 400

South Africa +27 11 254 8360

Spain +34 (91) 372 6055

Sweden +46 8 477 6503/4

Taiwan 886 (2) 2722-9622

United Kingdom & Eire +44 (0) 1344 392400

USA 1 (800) 426-2200

USA (Export Sales) 1 (503) 627-1916

For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

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 © 2002, Tektronix, Inc. 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.

07/02 RL/PP

55W-16025-0

