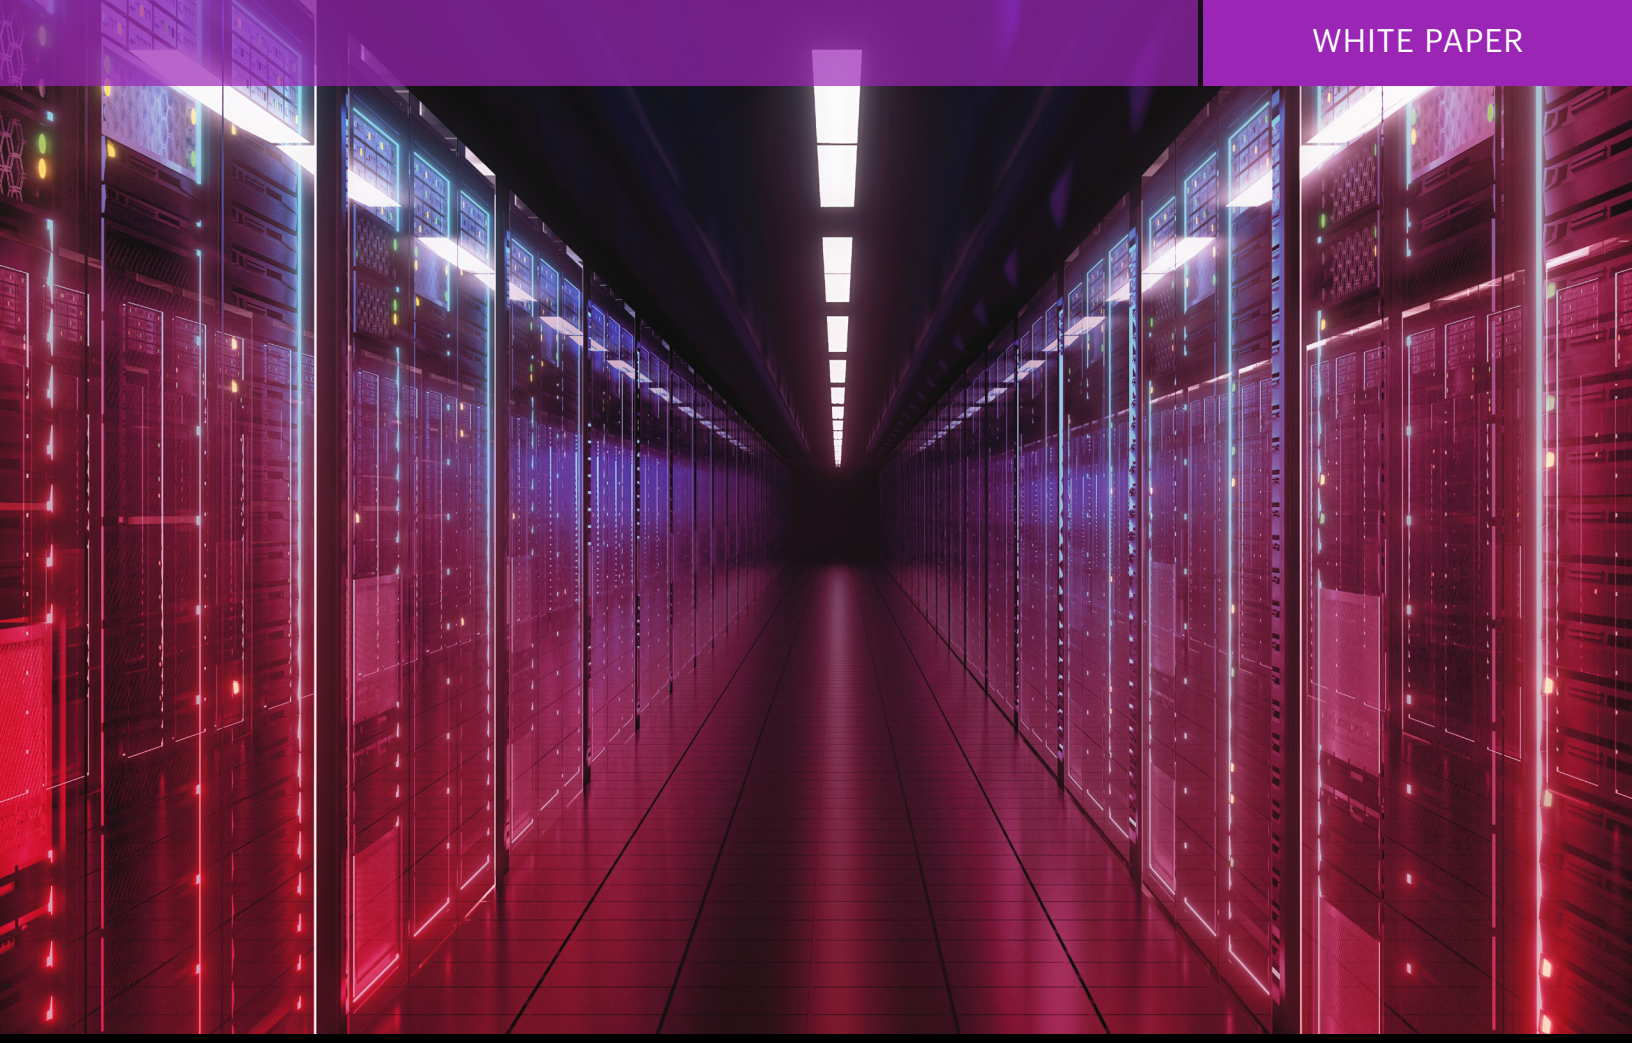
WHITE PAPER



Future-Proofing Networks with 1.6T Ethernet

Six interconnect challenges you must overcome

KEYSIGHT

Introduction

The rapid progress of AI places significant pressure on data centers, highlighting the critical importance of efficient scaling and optimization. Deploying 800GE and 1.6T Ethernet solutions is crucial in mitigating infrastructure bottlenecks. However, to ensure the reliable performance of these high-speed networks, we must validate high-speed interconnects and realistically emulate complex AI workloads using advanced test tools.

Engineers manually validated network interconnect performance for decades, a time-consuming process with limited or no automation that required advanced programming skills. This traditional approach lacks a centralized system for organizing and storing interconnect data and reports, making tracking and replicating tests and configurations difficult and time-consuming.

With the increased diversity and scale of AI and data center interconnects, traditional test methods can no longer accurately predict and measure the reliability of today's complex production networks. These methods struggle to keep pace with the rapid advancements and increasing complexity of modern AI and data center networks. Consequently, automated and AI-driven testing solutions are becoming increasingly prevalent in addressing these challenges.

This white paper explores the key challenges of validating interconnect performance in the context of 1.6T Ethernet. It then examines the increasing demands placed on data centers, the limitations of existing network architectures, and the emerging technologies and testing methodologies required to ensure the reliable deployment of these high-speed networks.

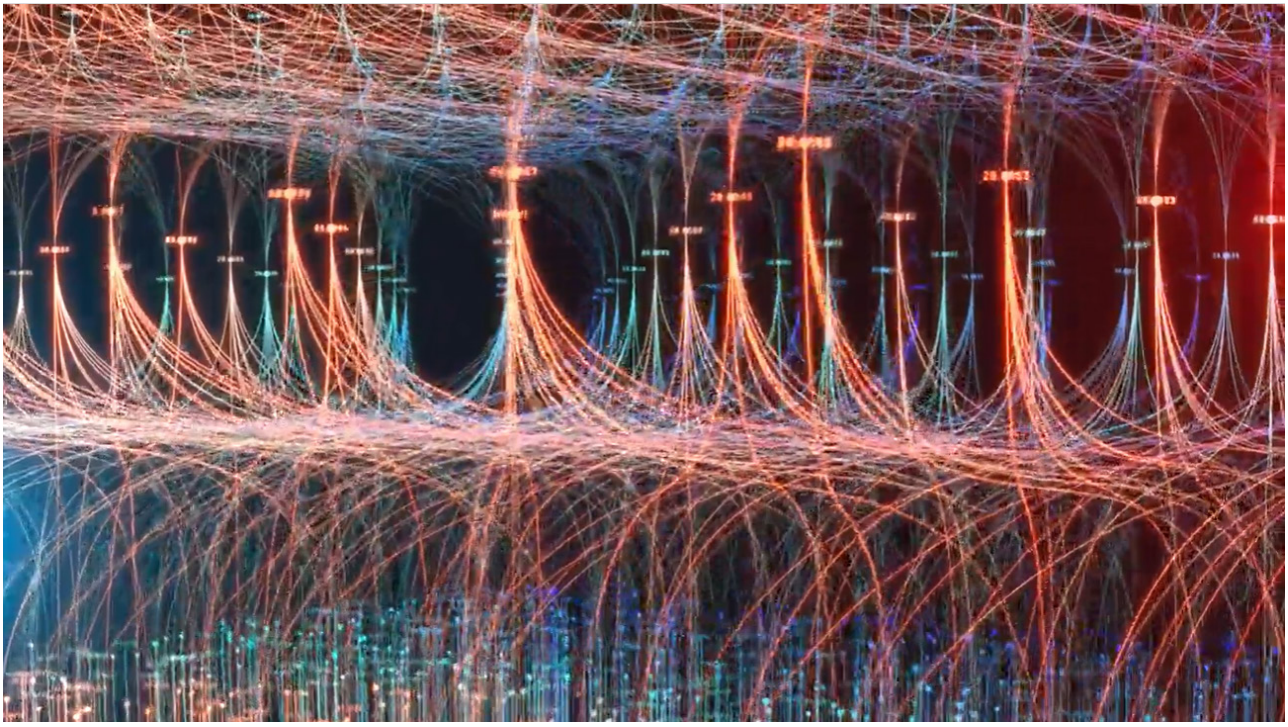
The paper also introduces the Keysight Interconnect and Network Performance Tester 1600GE (INPT-1600GE) and Keysight Interconnect Test System (ITS) software, which deliver a comprehensive, automated solution that streamlines interconnect validation and ensures the reliability and efficiency of high-speed network deployments.

The Escalating Demands of AI on Data Center Networks

The proliferation of AI applications drives an unprecedented surge in data traffic within data centers. AI workloads, such as deep learning, machine learning, and neural network training, are characterized by massive datasets and complex communication patterns. These workloads require the movement and processing of enormous amounts of information, often distributed across large clusters of interconnected computing resources.

Several key factors contribute to this escalating demand:

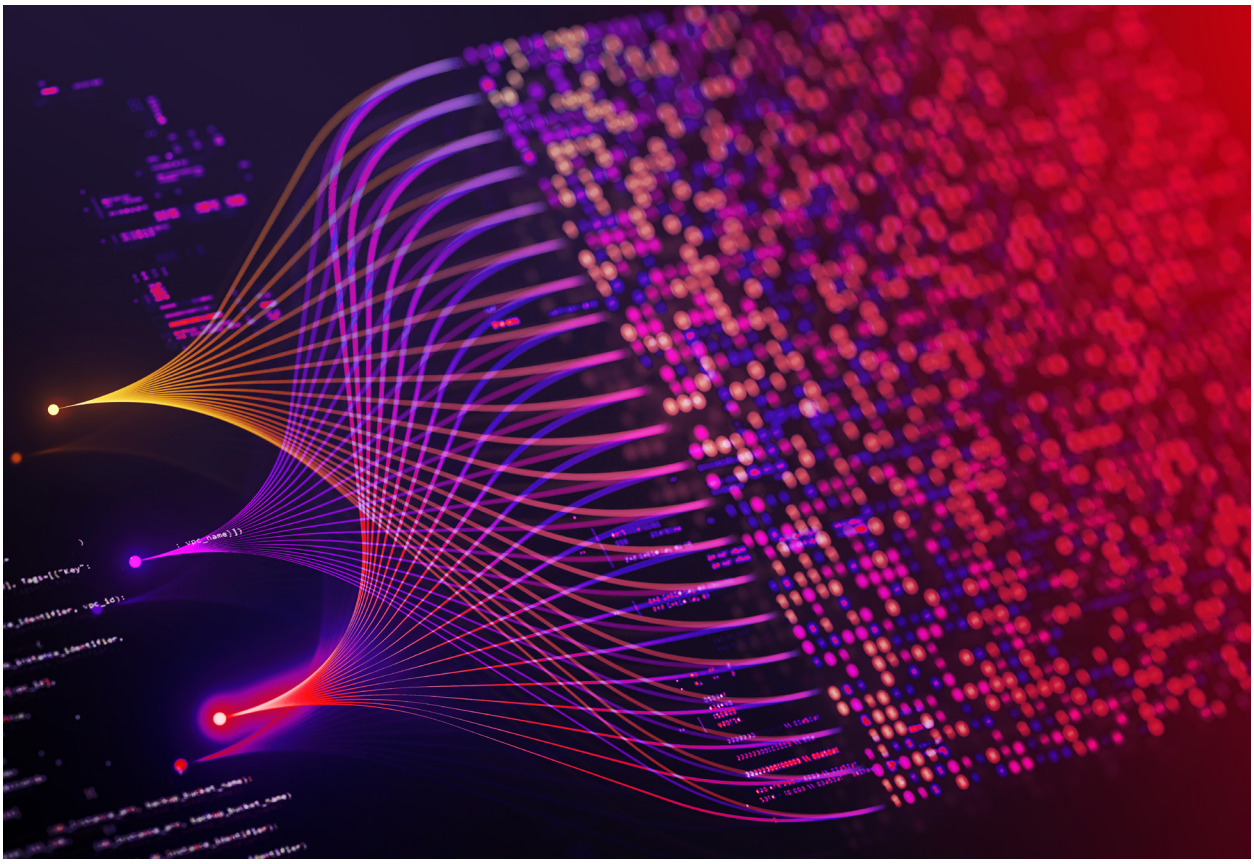
- **Exponential growth of data:** AI applications generate data at an exponential rate. This data needs to be stored, processed, and analyzed, placing immense strain on network infrastructure.
- **Increasing model complexity:** AI models are becoming increasingly complex, requiring more computational power and data to train effectively. This complexity translates to a greater demand for network bandwidth and throughput.
- **Distributed computing:** AI workloads are often distributed across multiple servers or GPUs, necessitating high-speed, low-latency communication between these resources.
- **Real-time processing:** Many AI applications, such as autonomous driving and natural language processing, require real-time data processing, placing stringent demands on network performance.



1.6T Is Here: Can Your Network Keep Up?

The networking industry is rapidly evolving to address these escalating demands of AI workloads. Advanced Ethernet technologies, such as 800GE and 1.6T, offer significantly increased bandwidth and reduced latency compared to previous generations. The exponential data growth driven by AI applications pushes the adoption of 800GE and 1.6T to ensure future scalability. These advancements enable networks to handle the massive data flows and complex communication patterns characteristic of AI applications.

1.6T Ethernet represents a significant leap forward in networking technology, offering a theoretical throughput of 1.6 terabits per second. This increased bandwidth is crucial for supporting the data-intensive workloads of AI and enabling the development of next-generation AI applications.



Limitations of Existing Network Architectures

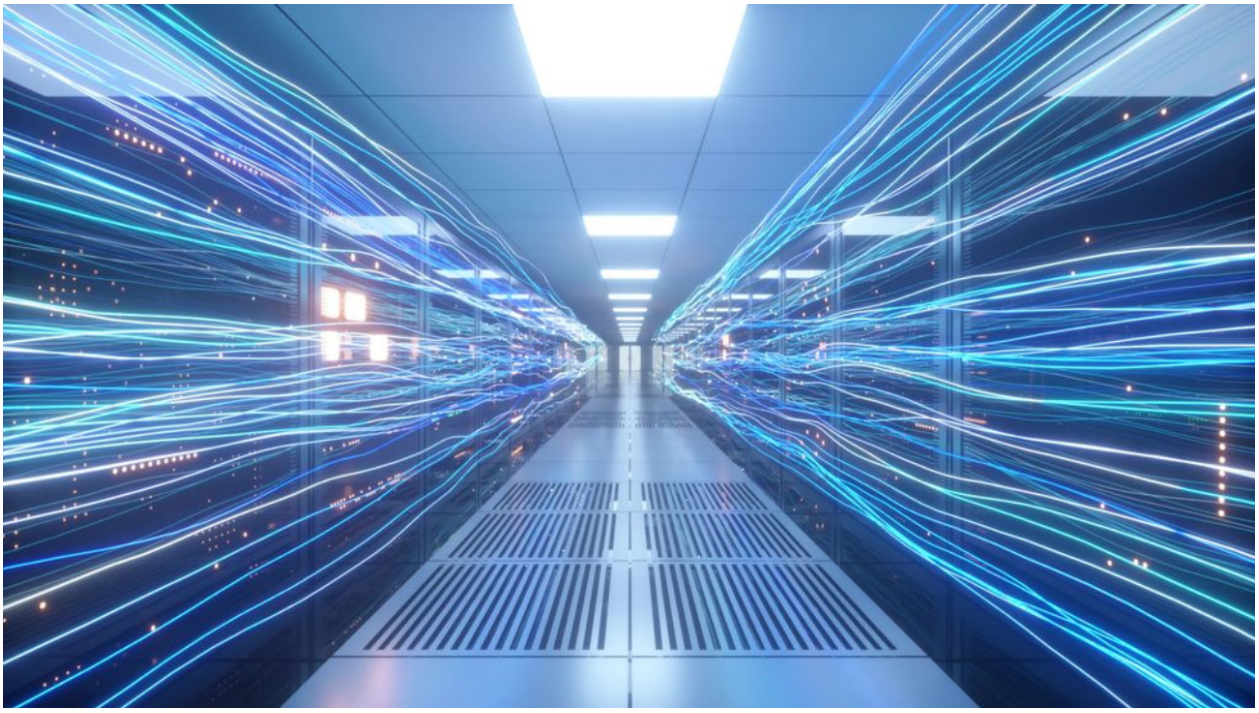
While 1.6T Ethernet promises transformative capabilities for AI infrastructure, most existing network architectures fall short of supporting such high-performance demands. Traditional data center networks were not designed with the scale, complexity, or bandwidth intensity of modern AI workloads in mind. As a result, they present critical limitations that hinder the full potential of advanced technologies like 1.6T Ethernet:

- **Latency bottlenecks:** AI workloads, particularly those involving distributed training, rely on ultra-low latency for efficient synchronization across nodes. Legacy networks introduce delays that reduce throughput and slow model convergence.
- **Bandwidth limitations:** Traditional Ethernet standards cannot match the bandwidth needed to move massive AI datasets quickly. This bottleneck causes data starvation and limits the speed of training and inference cycles.
- **Network congestion:** As data flows increase exponentially, oversubscribed links in conventional architectures lead to congestion, packet loss, and retransmissions — all of which reduce network efficiency.
- **Scalability challenges:** Scaling existing architectures to support thousands of GPUs or AI accelerators pushes the limits of switch radix, interconnect density, and power budgets.

These architectural constraints directly impact the effectiveness of AI deployments by causing:

- **Increased training times:** Delays in data movement elongate training durations, delaying time-to-insight and rising R&D costs.
- **Reduced model accuracy:** Packet loss and inconsistent data delivery introduce noise and instability in model training, potentially degrading accuracy.
- **Higher operational costs:** Inefficient data movement leads to higher power consumption, cooling demands, and operational overhead.

A shift toward modern, AI-optimized network architectures is imperative to fully realize the benefits of 1.6T Ethernet and meet the future demands of AI.



AI Workloads and the Interconnect Challenge

As AI workloads become more complex and data-intensive, they expose the limitations of traditional network infrastructures. Delays in data transfer create a major bottleneck for distributed AI workloads that require frequent, high-speed communication. Even minor delays significantly extend training time and reduce efficiency.

Limited data transfer capacity slows the movement of massive AI training datasets, potentially causing data starvation, where graphical processing units (GPUs) sit idle waiting for data. Excessive traffic saturates network links, worsening both latency and bandwidth limitations. Congestion causes packet loss, retransmissions, and delays, severely degrading AI performance.

These network challenges have profound implications for AI workloads and data center operations. Increased latency and reduced bandwidth slow down the training of AI models, leading to longer development cycles and higher costs. Network congestion causes unpredictable performance fluctuations, making it difficult to optimize AI applications.

1.6T Ethernet interconnects introduce six technical complexities

1. **Signal Integrity and PAM4 modulation**
   Testing PAM4 signaling presents a challenge due to its four-level modulation, which requires precise calibration to ensure signal integrity and minimize errors.

2. **Forward Error Correction (FEC)**
   Validating FEC involves complex algorithms to correct transmission errors, demanding rigorous testing to ensure reliable data communication.

3. **Common Management Interface Specification (CMIS)**
   CMIS testing is intricate as it involves standardizing the management and configuration of high-speed interconnects, requiring thorough validation to ensure interoperability.

4. **Linear Pluggable Optics (LPO)**
   Testing Linear Pluggable Optics (LPO) is difficult due to the need for modular and scalable deployment, which demands comprehensive validation to maintain network flexibility and performance.

5. **Bit Error Ratio (BER) and Frame Loss Ratio (FLR)**
   BER testing measures the rate of errors in transmitted bits, necessitating precise and continuous monitoring to ensure data accuracy. FLR testing quantifies the percentage of frames lost during transmission, requiring detailed analysis to assess network stability and performance.

6. **Optical Power Monitoring**
   Continuously monitoring optical power levels ensures optimal performance and prevents signal degradation.

Emerging Network Solutions for AI Workloads

Several emerging technologies are being developed and deployed to address the challenges posed by AI workloads:

- **Co-Packaged Optics (CPO):** CPO integrates optical devices directly with the switch ASIC, reducing the distance between the electrical and optical components. This integration offers significant benefits in terms of power consumption, signal integrity, and density, enabling the development of higher-speed and more efficient networks.
- **RDMA over Converged Ethernet (RoCEv2):** RoCEv2 is a network protocol that allows for direct memory access over Ethernet networks. RoCE improves the performance of AI workloads by enabling high-speed, low-latency data transfer between servers.
- **Time-Sensitive Networking (TSN):** TSN is a set of standards that enables deterministic, real-time communication over Ethernet networks. TSN is crucial for applications that require precise timing and synchronization, such as industrial automation and autonomous driving, and can also benefit AI applications that require coordinated data processing.
- **Intent-Based Networking (IBN):** IBN is a network management approach that uses high-level policies to automate network configuration and operation. IBN simplifies the management of complex AI networks and improves network agility and resilience.
- **Software-Defined Wide Area Networking (SD-WAN):** SD-WAN offers a flexible, software-driven approach to wide-area networking by connecting local area networks (LANs) across large distances using software that is compatible with various networking hardware. This technology enhances network flexibility, reduces operational costs, and improves cloud connectivity.
- **Wi-Fi 6 and 6E:** These standards represent the latest advancements in wireless connectivity, offering substantial improvements in speed, capacity, and overall performance. Wi-Fi 6 and 6E are particularly beneficial for environments with high device density and demanding applications.
- **Secure Access Service Edge (SASE):** SASE combines networking and security into a unified, cloud-native architecture. This approach addresses the challenges of managing diverse environments in a mobile-first, cloud-centric world, enhancing security and performance.
- **Microsegmentation:** This technology divides a network into smaller, isolated segments, making it more difficult for attackers to move laterally within an organization's IT infrastructure. Microsegmentation enhances security by enforcing granular access controls and tightly controlling traffic between segments.

These emerging technologies collectively contribute to the development of high-speed, efficient, and secure networks capable of meeting the demands of modern AI workloads.



Keysight's Philosophy on 1.6T Ethernet Interconnect Testing

Keysight believes that effective testing and validation are crucial at every stage in the development and deployment of 1.6T Ethernet solutions.

Scalable infrastructure is crucial for AI data centers to manage the increasing demands of AI workloads and evolving technology. Scaling AI data centers demands more than component-level validation. Interoperability, performance, and efficiency are system-level metrics that are only measurable under real-world network conditions. Therefore, test strategies must incorporate both physical-layer testing and full-stack AI workload emulation, which are capable of optimizing system-level performance.

Keysight emphasizes the importance of a comprehensive approach that combines hardware and software solutions to address these challenges. By focusing on scalable infrastructure and advanced testing methodologies, Keysight is committed to enabling its customers to achieve higher efficiency, reliability, and performance in their AI data centers, supporting the rapid advancement of AI technologies.

Five Advanced Testing Methodologies for 1.6T Ethernet

With the increased diversity and scale of AI and data center interconnects, traditional test methods no longer accurately predict and measure the reliability of today's complex production networks. Advanced testing methodologies are needed, including:

1. **Automated Test Case Generation and Execution:** Automating the creation and execution of test cases reduces testing time and improves test coverage significantly.
2. **Centralized Data Management and Reporting:** A centralized system for managing and reporting test data streamlines the validation process and improves collaboration.
3. **Realistic AI Workload Emulation:** Emulating real-world AI workloads in a test environment is crucial for accurately assessing the performance and reliability of network infrastructure.
4. **Comprehensive Layer 1 to Layer 3 Testing:** Assessing all layers of the network stack, from the physical layer to the application layer, is essential for ensuring interoperability and performance.
5. **Emphasis on System-Level Performance and Interoperability:** Evaluating the performance and interoperability of the entire network system, rather than just individual components, is crucial for ensuring reliable operation in a production environment.

Keysight addresses these challenges with the Keysight Interconnect and Network Performance Tester 1600GE (INPT-1600GE) and Keysight Interconnect Test System (ITS) software, as seen in Figure 1. The INPT-1600GE is a 1.6T Ethernet hardware test platform for traffic generation and analysis — and is part of the Keysight Artificial Intelligence (KAI) architecture, a portfolio of end-to-end solutions designed to help customers scale AI processing capacity in data centers.

Available in both a lightweight, portable benchtop unit and a 19-inch shelf-mounted rackmount chassis, the INPT-1600GE caters to diverse testing environments and needs. Working in conjunction with the hardware, the ITS software delivers a comprehensive solution that automates interconnect validation in high-speed Ethernet networks and AI data centers.



Figure 1. Keysight Interconnect and Network Performance Tester 1600GE shown with Interconnect Test System software

The ITS software, in Figure 2, includes a first-of-its-kind, US patent pending Interconnect Library (IL) that introduces a new way to organize, retrieve, and use interconnect data, including Common Management Interface Specification (CMIS) data. This cutting-edge library organizes all data and creates a comprehensive record that can be used to create and execute new tests.

Once a record is created, it is automatically added to the IL self-serve database, where users can retrieve, reuse, edit, and update records. This process dramatically increases productivity and accelerates the creation of automated test suites without advanced programming. The software features an intuitive browser-based GUI, enabling multiple users to run and schedule tests simultaneously, enhancing collaboration and productivity.

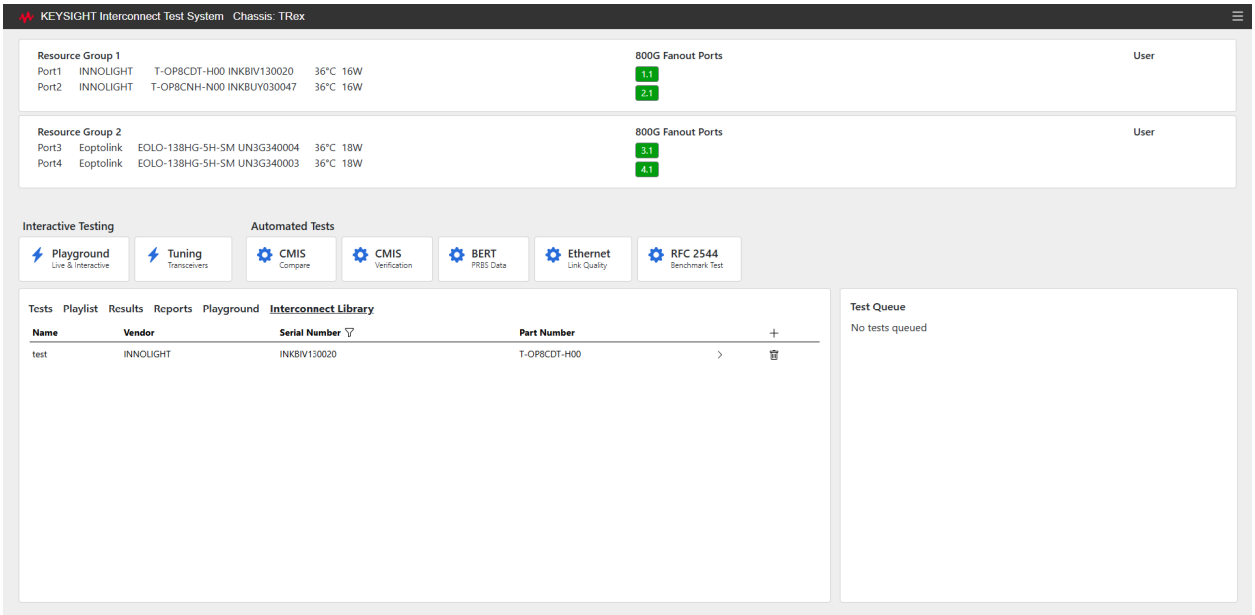


Figure 2. The intuitive home screen of Keysight's ITS software simplifies the process of configuring, running, and analyzing interconnect tests.

Conclusion

The increasing reliance on AI drives an unprecedented demand for high-speed, low-latency network solutions. Proactive and comprehensive testing is critical as AI workloads push network performance boundaries.

This white paper examines the escalating demands that AI workloads place on network infrastructure and the role of 1.6T Ethernet in addressing these demands. It also underscores the interconnect challenges of 1.6T Ethernet and the advanced testing methodologies required to ensure the successful deployment of this technology.

By embracing these advanced testing approaches, organizations can future-proof their networks, overcome the challenges of 1.6T Ethernet, and pave the way for next-generation, AI-optimized networks. Contact us to discover how our solutions can optimize your AI data center.

Learn More

Explore how Keysight solutions help validate high-speed Ethernet interconnects in AI environments:

- Product Page: Interconnect and Network Performance Tester 1600GE
- Solution Brief: Interconnect and Network Performance Tester 1600GE
- Flyer: Interconnect Test System Software
- White Paper: How AI / ML Networks Differ from Traditional Networks
- Blog: Key Challenges in Scaling AI Data Center Clusters
- Poster: Your Path to 1.6T
- Use Case: How to Validate Ethernet Interconnects in Data Centers

Keysight enables innovators to push the boundaries of engineering by quickly solving design, emulation, and test challenges to create the best product experiences. Start your innovation journey at www.keysight.com.

KEYSIGHT

This information is subject to change without notice. © Keysight Technologies, 2025, Published in USA, April 15, 2025, 7125-1022.EN