

Spirent Chip Design Verification Solution


Bridging the Gap Between Pre-Silicon Verification & Post-Silicon Validation

Application Overview

Growth in technology areas such as cloud computing, mobile edge computing, and 5G are pushing the demand for network capacity. Network equipment and semiconductor manufacturers need to keep up with the demand by delivering ultra-high-density devices powered by cutting-edge application-specific integrated circuit (ASIC) and system-on-a-chip (SoC) solutions.

Post-silicon validations are no longer sufficient to meet the aggressive time-to-market and cost-efficiency requirements. Design issues found late in post-silicon validation can cost thousands of hours of wasted effort and have a negative impact time-to-market. Access to scalable and accurate pre-silicon verification is becoming a must have for networking ASIC and SoC designers.

Spirent Chip Design Verification Solution is integrated and hosted together with EDA emulators. It is time-synced with EDA emulation clock for accurate and realistic Layer 2-3 traffic generation and real-time results analysis.



Spirent Chip Design Verification Solution enables effective and efficient pre-silicon validation from 1G to 400G, with 800G in the works. It bridges gaps between pre-silicon and post-silicon verifications, and delivers significant cost savings to customers by identifying issues early in the chip design stages. Using the unified Spirent TestCenter test platform in all phases of silicon product lifecycle provides:

- Improved operational efficiency and reduced total cost of ownership with zero learning curve
- More effective measurements and result analysis
- Reusability of test cases/scripts in pre-silicon verification and post-silicon validation

- **Cost savings**—identifying and addressing issues early in the chip design stages
- **Unified test platform**—bridges gaps between pre- & post-silicon verification, delivers significant benefits to users, including:
 - Capability to test all phases of silicon product lifecycle
 - Zero learning curve as the same application can be utilized
 - Use of standard metrics provides more effective measurements and result analysis
 - Reusability of test cases in pre-silicon verification and post-silicon validation
- **Accelerated time-to-market**—speeding up entire silicon development lifecycle



Features and Benefits

- Integrated with EDA emulators and time synced with emulation clock
- Support 1GE, 2.5GE, 5GE, 10GE, 25G, 40GE, 50GE, 100GE, 200GE, and 400GE emulated port speeds
- Accurate and realistic Layer 2-3 traffic generation
- Real-time results and powerful bandwidth and latency measurements based on emulation time enable efficient troubleshooting and fault isolation
- Custom Ethernet traffic generation with hundreds of predefined packet templates
- High-resolution latency measurements, with 2.5 nanosecond (ns) accuracy
- Automation capabilities with TCL and Python

Technical Specifications

Parameter	Description
Maximum # of Virtual Ports Per VM	8
VM Specifications	<ul style="list-style-type: none"> • CPU - 5 to 17 virtual CPU cores (1 vCPU core for control, the remaining vCPU cores are evenly spread over the virtual ports) • Memory - 2 to 4 GB
Packaging	Raw disk image, QCOW2
Hypervisors and Host OS	KVM over RHEL 7.3 / CentOS 7.3 (or newer) KVM over Ubuntu 16.04 LTS, and 18.04 LTS
vNIC Driver	Industry Standard Virtio Driver
Ethernet Port Speeds	1GE, 2.5GE, 5GE, 10GE, 25G, 40GE, 50GE, 100GE, 200GE, and 400GE
User reservation	Multi-user support. Per VM user reservation
Ethernet II Frame Size	58 – 9000 bytes
Traffic Generation	Raw traffic and L2/L3 bounded traffic
Source/Destination Ports Mapping	One-to-one, many-to-many, fully meshed
Scheduler mode support	<ul style="list-style-type: none"> • Port Based – traffic scheduling handled at the port level • Rate Based – key parameters determined at the port level with division among the individual stream blocks • Priority Based – scheduling determined at the stream block level using user-assigned priorities. Precise scheduling of CBR and bursty traffic for QoS testing
Latency measurement resolution	2.5 nanoseconds
Capture Buffer	2MB per port
Statistics per stream	Over 40 real-time measurements per stream—includes standard frame and packet counters and rates and advanced sequence checking, RFC 4689 jitter, latency, FCS errors and checksum errors. <ul style="list-style-type: none"> • Advanced sequencing: In-order, lost, reordered, late and duplicate • Latency: Avg, min, max and short-term avg; first/last frame arrival timestamp • Data integrity: IP checksum, TCP/UDP checksum, frame CRC, embedded CRC and PRBS bit errors
Statistics per port	Over 50 transmit stats per port reported in real time. Stats include Layer 1, Layer 2 and Layer 3+ counters and rates and include received FCS, checksum, and PRBS errors and rates.
Automation	TCL, Python

Ordering Information

Product Number	Description
AON-CD-V-SYN-016-SUB	AION Chip Design Verification Synopsys 16 Port Subscription
AON-CD-V-SYN-064-SUB	AION Chip Design Verification Synopsys 64 Port Subscription
AON-CD-V-SYN-256-SUB	AION Chip Design Verification Synopsys 256 Port Subscription