PRODUCT

Network Switch

Telecommunications Network Switch Zynq Ultrascale Mpsoc

SiliconCentric | Independent Verification & Validation

SiliconCentric

I

PRODUCT

Introduction

Our client is a telecommunications equipment vendor. Their product is a 12-port 10 Gb aggregation switch, with four first-article boards back from fabrication and a carrier-customer field trial eight months out. The design is a standalone Zynq UltraScale+ MPSoC — packet datapath in the programmable logic, management plane on the Cortex-A53 cores, and the internal RPU acting as the companion controller for board management, fan control, and PSU supervision. We collaborated with them on this project by providing independent verification and validation of the board.

- Compute topology: a single MPSoC — fabric datapath, A53 management plane, and the internal Cortex-R5 RPU handling board-management duties a discrete BMC would own on a larger chassis.
- Under validation: 12 SerDes lanes at 10.3125 Gbps into SFP+ cages, two multi-rate PHYs, the SFP+ I2C management tree behind a mux, PTP/1588 timestamping hardware, a PMBus-instrumented dual PSU stage, DDR4, QSPI boot flash, and the fan controller.
- Program state: first articles fabricated, powered for basic bring-up only, management-firmware team scheduled and waiting.
- Independence carried weight with the carrier customer: an assessment with no technical, managerial, or financial stake in the design.
- Scope was bounded to independent validation — no feature development, no switching-stack work, no firmware authoring for the product.
- We worked from the hardware design package alone: schematics, ICDs, BOM — no access to firmware source.

2

PRODUCT

Hardware Topology

This is a standalone-SoC topology with the companion function on-die: the RPU is the board manager, the A53 cluster is the management plane, the fabric is the datapath. Scope enclosed the physical board — lanes, cages, clocks, power, thermals; the switching stack and any carrier acceptance testing sat outside.

- Topology classification: standalone MPSoC with internal RPU companion; APU–RPU interconnect via shared OCM and IPI mailboxes, with the RPU owning fan PWM, PSU PMBus, and thermal-trip authority.
- The RPU's job: fan-curve control, PSU telemetry and failover, temperature supervision with hard trip authority, and power sequencing of the PHYs and optics rails.
- In-scope interfaces: SerDes x12 (10GBASE-R), PHY management (MDIO), SFP+ I2C behind an 8-channel mux, PM-Bus to both PSUs, DDR4, QSPI, RS-232 console, out-of-band GbE management port, PTP hardware timestamp path.
- Envelope in scope: 0 °C to +55 °C inlet ambient per the telecom-shelf spec, dual-feed –48 V DC input, and continuous full-traffic duty.
- Out of scope: the switching stack, carrier field-trial acceptance, NEBS formal testing — we generate the pre-compliance and integration evidence; the client's lab and their carrier judge it.
- Zynq UltraScale+ is among the six architectures Polaris already runs on, so the OS baseline was inherited and engineering went to this board's delta rather than silicon bring-up.

3

PRODUCT

Problem Statement

Four boards existed, a field trial with a paying carrier was on the calendar, and the vendor's credibility rode on those twelve ports staying up. Every SerDes, PHY, and PSU had vendor characterization behind it; the board connecting them had none.

- Proven in isolation, never as a board: Xilinx characterized the GTH transceivers on their own test silicon and the PHY vendor on an evaluation board — but this board's twelve-lane breakout routing, its cage stack, its clock distribution, and its dual-PSU sequencing had never been exercised together.
- Firmware coverage is feature-shaped, not board-shaped: the management firmware would drive the datapath and the console daily, while the SFP+ mux tree's fault modes, the second PSU's failover path, PTP under load, and the fan-controller edge cases fell outside every planned feature.
- No stress, no extremes, no duration: no 55 °C inlet soak at full traffic, no multi-day saturation endurance, no dual-feed failover and brown-out campaign, no lane-margin characterization across temperature — the operating reality of a telecom shelf.
- Nothing runs alongside development: validation was a pre-trial gate, so a marginal lane would first appear as CRC errors in the trial — debugged remotely, on a carrier's site, at the worst possible time.
- The carrier's acceptance regime expects error-free soak evidence and PTP conformance data; NEBS thermal assumptions expect the board to survive its own shelf.
- 10.3125 Gbps lanes depend entirely on board-level signal integrity — eye margin is a property of this routing, this stackup, this cage, at this temperature.
- A field-trial failure costs the vendor the carrier relationship; a fleet recall after deployment costs multiples of the entire validation budget.

PRODUCT

Customer Requirements

4

The client's requirements were direct: validate the whole board, under stated limits, with evidence shaped for the carrier's acceptance reviewers.

- Operating temperature: the board shall operate from 0 °C to +55 °C inlet ambient at full traffic.
- Supply voltage: dual-feed −48 V DC telecom input, operating window −40 V to −57 V. The board shall survive input to −60 V without damage; reverse-polarity and over-voltage protection shall be verified at the power entry.
- Feed failover: loss of either feed shall cause no packet loss and no reset; the board shall ride through feed brown-out to −36 V for 10 ms.
- Over-current: each PSU output and each optics rail shall have over-current protection tripping within 10 ms with PM-Bus fault reporting; a shorted SFP+ module shall not disturb any other cage.
- Thermal protection: the fan controller shall detect a stalled fan via tach fault; the RPU shall hold hard thermal-trip authority. The board shall not reach thermal shutdown at 55 °C inlet with all fans running.
- Humidity: the board shall operate at up to 85% RH non-condensing; formal NEBS environmental testing sits with the client's lab.
- Full-board functional validation: all twelve lanes at rate, both PHYs, every SFP+ cage and its management channel, PTP timestamping, both PSUs with failover, DDR4, QSPI, console, management port, fans — one integrated system.
- Coverage independent of firmware scope: every subsystem validated regardless of the management-firmware feature set, through a HAL built from the design package.
- Stress matrix: 0–55 °C inlet sweep with lane-margin scans at each step, 168-hour full-traffic endurance, dual-feed failover and −48 V brown-out campaigns, bus saturation with PTP conformance measurement — vibration deferred to the NEBS lab.
- Continuous validation alongside development: CI integration revalidating every management-firmware candidate on real boards.
- Evidence: timestamped, operator-attributed, version-pinned records with bidirectional traceability, packaged for carrier acceptance review.
- A web console the vendor's lab technicians run without knowing the validation stack — learned once, identical across programs.
- An immutable, append-only evidence store carrying first-article baselines forward to the production run.
- Setup in minutes per board: standardized procedure, no custom rig, repeatable across all four first articles.

PRODUCT

Architecture

The board runs the validation OS, HAL, and client; an external x86 host runs the server, dashboard, and SDK. Test execution stays isolated from data processing and the MPSoC keeps its resources for representative behavior — necessary here because lane margins had to be scanned while the fabric datapath ran at full rate, not on an idle board.

- Polaris OS (the validation OS) — a reproducible Linux image for the A53 side booting to a known state, one baseline across four boards instead of four bench setups.
- HAL — the one board-specific layer, exercising the GTH lanes and eye-scan hardware, MDIO, the SFP+ I2C mux tree, PMBus, PTP timestamp path, DDR4, and fan control behind a uniform interface that keeps the telecom suites portable to the production revision.
- Client — the on-target agent driving tests through the HAL; gRPC on Linux, FlatBuffers to the RPU's FreeRTOS image; three timestamps per event (origin, relay, server); light enough that PTP measurements reflect the shelf, not the tooling.
- Server — orchestration, scheduling, and the append-only records store on the host, four boards in parallel, dashboard and API included.
- Dashboard — the lab console: provision a board, configure and launch, live per-lane error counters, eye-margin plots, PSU and thermal telemetry, time-series review; every state-changing action operator-attributed.
- Python SDK — automation and CI; scripts the console's full capability; C bindings available, unused here.
- The OS and HAL are the two board-specific layers and were heavily modified for this board and this SoC; the client, server, dashboard, and Python SDK were configured to the client's requirements, but their core structure and working style remain the same across every board and every industry.

6

PRODUCT

Design

The binding constraints: twelve lanes whose margins shrink with temperature, a 55 °C inlet spec, dual-feed telecom power, and a soak-evidence culture where carriers ask for the raw counters.

- Validation OS design — read-only rootfs, A/B partitions with rollback, boot from QSPI with fallback image, because re-flashing a board mid-thermal-sweep forfeits the sweep; the trade-off is QSPI image-size discipline, enforced. The RPU runs FreeRTOS, matching the vendor's eventual BMC-function firmware environment so findings transfer.
- HAL design — partitioned by subsystem against schematic sheets and ICD clauses; lane statistics and PTP events are interrupt-driven while optics telemetry polls at 1 Hz through the mux — the polling choice deliberate because driving the I2C mux faster than the optics' EEPROM timing allows produces phantom NACK findings, a pitfall we documented after reproducing it.
- Client design — origin timestamps taken from the 1588 hardware clock so PTP evidence is self-consistent; under thermal-trip the client flushes records before the RPU cuts rails, trading shutdown speed for a complete record of every trip event.
- Server design — three-timestamp causality orders a PSU sag against the lane errors it caused; retention sized for 168-hour runs of per-lane counters on four boards, with raw eye-scan matrices stored compressed — bulky, but the carrier wanted originals.
- Dashboard design — per-lane BER counters, eye-margin heatmaps, PSU rail telemetry, and inlet/die temperatures are first-class views; run health and drop counts on every screen, because a lab tech must spot an evidence gap before a carrier reviewer does.
- SDK design — the vendor scripted the acceptance gate, the thermal-margin matrix, and the failover campaign; the direct API path served single-lane bench checks, at the stated cost that such runs produce no records.
- OS and HAL design was board-specific and heavily reworked; the four upper components were configured to the vendor's lab workflow while their structure and working style stayed constant.

7

PRODUCT

Implementation

Project work went to the OS delta, the HAL, and the telecom campaigns; the client, server, dashboard, and SDK required configuration only.

- Validation OS — Yocto image for the MPSoC: inherited boot chain (FSBL, ATF, U-Boot), board device tree covering PHYs, mux tree, PMBus devices, and fans; minimal kernel; OTA with A/B rollback; boot to validation-ready in under 30 seconds; per-build SBOM with CVE tracking; RPU FreeRTOS image in the same tree.
- HAL — built subsystem by subsystem from the design package, bench-verified as written: register interfaces to the GTH eye-scan engines, MDIO drivers, mux-aware I2C paths with per-channel timing, PMBus drivers for both PSUs, DMA capture of PTP event timestamps; the package pass surfaced a schematic-to-BOM mismatch — the SFP+ management mux's reset line showed a pull-up on the schematic absent from the BOM, leaving reset floating and explaining an intermittent mux lockup seen in bring-up.
- Client — C++ daemon, event and polling paths, hardware-clock timestamping, graceful shutdown without record loss, gRPC upstream, FlatBuffers to the RPU.
- Server — telemetry receiver, versioned schema, scheduler, REST API, dashboard hosting, four-board parallelism.
- Dashboard — board provisioning, campaign configuration, live lane and thermal streaming, eye-map rendering, time-series review, export for carrier review, operator attribution.
- SDK — Python bindings, per-subsystem example scripts, and CI wiring firing the acceptance suite on every management-firmware candidate.
- Board-specific: OS delta and HAL. Configuration: client transports, server retention and compression, dashboard views, SDK examples.

PRODUCT

Test - IV&V

Testing proceeded stack-first, then OS/HAL, then the board through one-shot, long-running, event-based, and monitoring campaigns — all dashboard-launched runs committed to the append-only store.

- Stack self-verification: 100% test coverage of the client, server, dashboard, and Python SDK before it produced evidence about the switch.
- OS/HAL verification: A/B boot and rollback exercised, OTA verified end to end, secure boot checked on accept and reject paths.
- Bench setup: board cabled to the host, loopback optics and traffic generator connected, programmable −48 V dual-feed supply configured, provisioning from the dashboard — under thirty minutes per board.
- Functional campaign: per-lane BER suites at 10.3125 Gbps, eye-scan margin runs, PHY register and MDIO suites, SFP+ mux-tree and optics-EEPROM suites, PTP conformance under load, PMBus telemetry and PSU failover, DDR4 margin, console and management-port checks.
- Compute and memory checks: A53 cache and DMA integrity exercised under full datapath saturation, management-plane CPU load driven to worst case while lane statistics stayed clean, and QSPI boot-flash margin verified across the inlet range.
- Watchdog and reset campaign: induced A53 hangs confirming RPU supervision and recovery with the fabric datapath forwarding undisturbed, warm and cold reset combinations verified against the APU–RPU boot-order contract, and the RPU's thermal-trip authority fault-injected independently of the (reworked) fan-fault path.
- Debug posture: JTAG lockdown verified for the deployment configuration and console access checked against the management ICD — a carrier shelf exposes exactly the interfaces its spec says it does — with debug retained on first articles for eye-scan instrumentation.
- One-shot tests as the acceptance gate: one pass/fail per subsystem per firmware candidate per board.
- Long-running: the 168-hour full-traffic soak with periodic eye snapshots produced the main finding — lane 7's eye closing below mask margin above 48 °C inlet, traced to a via-stub resonance on that lane's breakout, corrected with back-drilling in the production layout.
- Event-based: silent runs recording on BER thresholds, PTP offset excursions, rail events, and thermal crossings — this mode caught PTP offset spikes correlated with PSU load-sharing oscillation between the two feeds, a firmware-tunable PMBus setting found before the trial.
- Monitoring: passive telemetry — inlet and die temperatures, all rails, fan RPM, per-lane counters — under sustained full traffic; a second serious finding surfaced here: the fan controller's tach-fault detection never fired when a fan was stalled by hand, and the board ran to thermal shutdown at 55 °C inlet — the RPU trip saved it, the fan fault logic did not, and the detection circuit was reworked.
- Instrumentation correlation: high-bandwidth scope and BERT captures anchored the lane-7 and PSU findings to the records at the electrical level.
- Recording rule: every run launched from the dashboard is recorded, always, into the append-only store; the Python SDK can run 100% of what the dashboard can run — module, peripheral, or system level — for quick independent checks, and SDK-launched runs of this kind are not written to the records store. The store is the evidence; the SDK is the bench tool.
- Traceability: requirements traced bidirectionally to tests and records, schemas version-pinned, runs reproducible for the carrier's reviewers.

9

PRODUCT

Deliverables

The vendor received the complete platform fitted to this board — source, records, procedures, no lock-in.

- Validation OS image with the full Yocto build system, recipes, and per-build SBOM.
- Board-specific BSP and HAL source for every subsystem, maintained against the design package into the production layout.
- Compiled client for the A53 side plus the RPU FreeRTOS build.
- Server and dashboard as deployment-ready containers in the vendor's lab.
- Complete suites — lane margin, BER, PTP conformance, PSU failover, thermal, 168-hour endurance — with recorded first-article baselines including raw eye-scan data.
- Python SDK with documentation, example scripts, and CI integration gating every firmware candidate.
- The append-only records store as vendor property — the evidence package for carrier acceptance.
- The independent validation report: lane-7 eye closure, the fan-fault detection hole, the mux reset pull-up mismatch, the PSU load-sharing interaction, metrics, pass/fail, coverage.
- Lab procedures for board acceptance and re-runs; full documentation; no lock-in.

10

PRODUCT

Conclusion

The board entered its carrier field trial with its margins measured, its worst lane fixed, and its thermal-protection story verified — the trial-killing defects had already been found on the bench.

- Lane 7's temperature-dependent eye closure and the silent fan-fault hole were both invisible at room temperature and both fatal in a hot shelf; catching them cost a back-drill note and a circuit rework instead of a failed trial and a damaged carrier relationship, with raw-counter evidence the carrier's reviewers accepted directly.
- The suites now gate every management-firmware drop in CI and carry to the production revision with only the HAL delta re-verified. The standing caveat holds — we produce the evidence; the client's lab and carrier judge it.

II

PRODUCT

Lessons Learned

This engagement made temperature-swept lane characterization a fixed early milestone on any multi-gigabit board we validate.

- Eye margin at room temperature is close to meaningless as evidence on 10 Gb-class routing — lane 7 passed every 25 °C scan and failed at 48 °C, so margin-versus-temperature sweeps now run before, not after, functional depth on any SerDes board, and fault-detection circuits (fan tach, PSU alerts) get physically injected faults rather than register-level simulation.
- The RPU-as-BMC pattern has a boot-window trap: for roughly 900 ms after power-on the fans sat at the hardware-default duty while the RPU image loaded from QSPI, which at 55 °C inlet is a real thermal excursion — on future MPSoC board-management designs we validate the pre-firmware power and cooling state as its own campaign.

CONTACT

SiliconCentric

EMAIL sandesh@siliconcentric.com

PHONE (669) 356-1998

WEB https://www.siliconcentric.com

ADDRESS San Jose, CA, USA