


Why test failures caused by USB infrastructure are often mistaken for code defects

By [Nicholas Metcalfe](#), Head of Technical Sales Engineering, Cambrionix

When a test run fails at 2 a.m., everyone looks at the code first. But in large-scale device labs, the culprit is often invisible, and it's sitting on the bench between the host computer and the device.

There's a certain kind of failure that doesn't show up cleanly in a test report. The device dropped its USB connection mid-run. The suite was interrupted. The result was logged as failed, but not because of anything in the build. A rerun conducted hours later passes without a change to a single line of code.



This isn't a fringe scenario. Across app testing pipelines, OS validation labs and managed cloud device farms, infrastructure instability is a common and one of the least visible causes of failed test runs. And it's one of the most expensive in terms of engineering time, delayed releases and eroded trust in the test suite itself.

Tests fail for two kinds of reasons. Only one of them is your code

Software test failures are expected to be deterministic: the same inputs, the same conditions, the same result. When a test is inconsistent - passing sometimes and failing others without any code change - it undermines the value of the suite. Engineers start ignoring red results and releases get delayed while teams manually distinguish real failures from infrastructure noise.

In real-device testing environments, a device can fail a test for reasons entirely unrelated to software: a dropped USB connection mid-run, a device that ran out of charge, a hub that couldn't sustain stable bandwidth under load, or a reboot triggered by a power interruption rather than a test instruction. These are infrastructure failures. They don't show up in your diff. They don't correlate with your commits. And in high-density labs, they can affect hundreds of devices simultaneously.

What happens when a USB hub can't keep up

Consumer-grade and legacy USB 2 and USB 3 hubs were not designed for sustained, high-density workloads. In a real-device testing environment, the demands are fundamentally different:

- **Contention on shared buses** causes dropouts when multiple devices transfer data simultaneously
- **Insufficient power delivery** means devices run down during heavy test cycles and tests abort
- **Transient load spikes** from simultaneous reboots can fault the hub, dropping all connected devices at once
- **No per-port observability** means engineers discover failures hours later with no information about what went wrong

Each of these produces test results that are corrupted, incomplete or simply wrong - and none of them are caused by the code under test.

The three teams paying the highest cost

App testing teams run suites triggered at each commit. A connection drop mid-suite logs as a code failure. The engineer investigates, finds nothing, reruns and it passes. An hour spent on a ghost. The CI/CD model depends on deterministic signal. Infrastructure noise destroys it.

OS and platform teams operate at a different scale. A single cycle may involve coordinated firmware flashing across hundreds of devices, multi-stage reboots, mixed OS versions. These runs take hours. One hub instability - one port losing charge mid-flash - can require a full restart. Time loss is measured in engineering-days.

Cloud and managed device farm operators face a different pressure: their infrastructure reliability is their product. When their farm has issues, their customers have issues. SLAs are tied to uptime. A failure here carries reputational and contractual consequences.

Across all three, the most damaging long-term effect is subtler: engineers learn that some failures are noise, start discounting red results, and stop trusting the suite. The investment in automation erodes, not because the tests are bad, but because the infrastructure made them unreliable.

What good looks like

A well-run device lab has one defining characteristic: engineers almost never investigate a failed test that turns out not to be a code failure. That requires infrastructure that sustains full bandwidth to all ports under load; delivers sufficient, consistent power during heavy workloads; surfaces problems immediately with per-port precision; and isolates faults so one device issue doesn't cascade across the hub.

If your suite has unexplained flakiness, devices losing charge mid-cycle, or engineers chasing failures that turn out not to be code issues - audit the USB infrastructure first. The fix doesn't make testing faster by itself. It makes it reliable. And reliability - the certainty that a failure means something real - is what makes automated testing worth the investment.

The best test runs look uneventful. That's not a low bar. It's the result of infrastructure built to handle what testing actually demands.

