FIBER OPTIC FIELD GUIDE

Troubleshooting a distributed-split network

Rural RDOF builds keep landing on the same design: a distributed split with the splitters fusion-spliced into the plant. It saves fiber, and it takes away the easy way to test. You can't break in at a splitter cabinet and shoot each leg, because there is no cabinet. Here is the field method that actually works, and why the wavelength you pick depends on which half of the network you're standing in.

Johnstone Technical Services · https://johnstonetechs.com/guides/distributed-split-ftth-troubleshooting

A centralized split puts all the splitters in one connectorized cabinet. When something breaks, you open the cabinet, disconnect a leg, and OTDR it dark. A distributed split spreads the split across the plant. GPON runs a 1:64 split, and we see it reached a few ways: straight 1:64, or distributed across two stages, a 1:2 at the feeder into 1:32s, or a 1:4 into 1:16s. In the builds we see most, those splitters are fusion-spliced into handholes, pedestals, and aerial spans. There is no leg to disconnect. The only places a tech can physically connect are the feeder at the OLT, or a live subscriber drop. And the fiber is carrying service the whole time.

So the instinct to shoot one clean trace through the whole PON is the wrong instinct. Past the splitter, every branch leg reflects back on top of every other, the cumulative splitter loss buries the small events, and you can't take the plant dark without knocking out subscribers. The method that works does the opposite: it splits the network at the splitter, tests each side from the end you can reach, and matches the wavelength to whether that side is live or dark.

Segment 1: the feeder, OLT to splitter

This is the stretch from the OLT PON port out to the first splitter. To test it, you disconnect the fiber at the OLT, which pulls it out of service and makes it dark. Dark is good: you get the clean, high-resolution trace that 1550 nm gives you, with none of the complication of live traffic on the glass.

The part that trips people up is worrying about the outage. Here is the thing to internalize: if the fault is on the feeder, every subscriber on that PON is already down, because the feeder is upstream of the split. Pulling the fiber to test it does not create an outage that didn't already exist. You are free to take it dark and shoot it forward.

Segment 2: the last mile, drop to splitter

This is the stretch from a subscriber drop back to the splitter. Here you usually can't go dark, because the fault is on one drop while the other subscribers on the PON are still in service. Take this segment dark and you knock out customers who were working fine.

So you test it live, and you test it from the subscriber end. Connect at the NID, and shoot reverse, back toward the splitter, using an out-of-band wavelength that doesn't touch the traffic. That wavelength is 1650 nm, filtered, which is the whole point of a live-fiber OTDR.

Free in your browser: OTDR trace analyzer, splice matrix builder, fiber analysis, link loss calculator, and more at johnstonetechs.com.
Informational guidance; verify critical measurements against certified test equipment.

Why 1550 one way and 1650 the other

The wavelength is not a preference, it is set by whether the fiber is carrying traffic. PON traffic lives roughly between 1270 and 1577 nm, with the downstream and any RF video overlay sitting right in the band an OTDR would normally use. On a dark feeder that doesn't matter, so you use 1550 nm and get the cleanest trace. On a live last mile it matters a lot: shoot 1310 or 1550 into live fiber and you step on the downstream.

That is what 1650 nm is for. It sits above the whole PON band, including XGS-PON at 1577 and the newer NG-PON2 downstream near 1600, far enough out that a filter in the OTDR passes your test pulse and blocks the live traffic from the receiver. It is the ITU-designated access maintenance wavelength, and the blocking filters built into ONTs are designed around it. If you also do transport or long-haul work, you'll run into 1625 nm, which is the equivalent out-of-band wavelength on C-band systems and what those acceptance specs call for. The two are not interchangeable: 1650 for access, 1625 for transport.

What you're hunting: macrobends, not splices

In an active network you are usually not looking for a bad splice. Those get caught at turn-up, and 1310 nm is where they show up. What takes down a network that was already working is physical damage, and damage almost always lands as a macrobend: a tree comes down on an aerial span, soil shifts on a buried run, or someone re-enters a splice point and pinches a fiber putting it back.

Macrobends give themselves away at 1550. A bend leaks more light the longer the wavelength, so you shoot both and compare: loss that is small at 1310 and bigger at 1550 is a bend, not a splice or a connector. That comparison is the main reason the dark feeder shot runs 1550, with 1310 as the confirming second shot. It is also why real dynamic range matters, because a bend's loss can be subtle until you've pushed past the splitter budget to see it.

Location is everything. A macrobend on the feeder, before the splitter, drags down every subscriber behind it, so one pinched fiber in a handhole can read like a whole-PON outage. That is exactly why you check the feeder first, and why you can take it dark to test without a second thought.

The hard case: a bend hiding at the splitter

The worst version is a macrobend right at the splitter, on the input side. Shooting forward, the tech sees the splitter on the trace, a big loss step, with the output legs past it. The trap is that a splitter is supposed to lose a lot, around 20 dB on a 1:64, so a big step looks normal. It is only a fault if you know it is bigger than it should be.

That takes two skills, and it is where outages get long. First you have to measure the loss, which means placing the analyzer's A/B markers across the splitter event correctly, not eyeballing it. Then you have to know what to expect, which is the splitter's rated loss plus a nominal splice. The Link Loss Budget Calculator gives you that number. If the measured loss comes in meaningfully above the budget, the extra is not the splitter. It is a bend sitting on the input.

And it only shows if you are on the right wavelength. A bend at the splitter adds more loss at 1550 than at 1310, so the event reads worse at 1550 by more than the design allows. Shoot only 1310, or never compare

the two, and the bend hides inside a splitter loss you already decided was normal. I have watched techs burn half a shift on an outage, staring at a splitter they could see on the trace the whole time, blind to a macrobend on the input, because they did not know what loss to expect and were not on the wavelength that would reveal it.

Put the two traces together

Neither shot alone finds the fault. The feeder trace tells you whether everything from the OLT to the splitter is clean. The last-mile trace tells you whether the drop back to the splitter is clean. Read them together and the fault falls into one segment or the other, with a distance inside that segment. If the feeder is clean and the drop shows an event at 1.2 km, you know where to dig. If the feeder shows the problem, you also know, before you ever roll to a drop, that this is a full-PON outage and not a single-subscriber ticket.

The splitter is a wall, not a window. A 1:64 split, straight or distributed as 1:2 into 1:32 or 1:4 into 1:16, stacks roughly 20 to 21 dB of splitter loss, before fiber and connectors. That is why you test each side from its own end instead of trying to shoot through, and why a live-fiber OTDR for this work needs real dynamic range at 1650, not the 24 to 28 dB a bargain fault locator gives you. Shooting through a cascaded split is how you end up staring at a noise floor.

Know the service impact before you pull

The method carries a nice operational side effect: it tells you what you're touching before you touch it. A feeder fault is already a full-PON outage, so taking that segment dark to test costs nothing. A drop fault is one subscriber, so that segment stays live and you test it out-of-band. Deciding which case you're in, before you disconnect anything, is the difference between a clean troubleshooting call and an outage you caused yourself.

The hard part is reading the trace

On a distributed-split plant the trace is rarely a clean slope with one obvious event. You're accounting for splitter loss at each stage, comparing measured loss against what the design should produce, and separating a real fault from the normal loss of the split. That interpretation is where the time goes, and where mistakes get made.

A few of our free tools take the arithmetic off your plate. The Link Loss Budget Calculator builds the expected budget, splitters and all, so you know what a clean 1:64 should measure before you ever shoot it. Because a distributed split is really a splice job, the Splice Matrix Builder documents the whole design, splitters included, so the plan you test against lives in one place. And once you have a trace, the Fiber Analysis tool reads it and does the event and loss work for you, with the same engine available to your own acceptance stack through the JTS API. The method above gets you a good trace from the right end of the network. The tools turn it into an answer.