CUTI FTTH DEPLOYMENT & MANAGEMENT — WEEK 3

M3.2 Overview of Splicing & Optical Validation

Delivery: hybrid — 6 h facilitated + 6 h self-paced · **Prerequisites:** M1.1, M1.3 · **Version:** v1

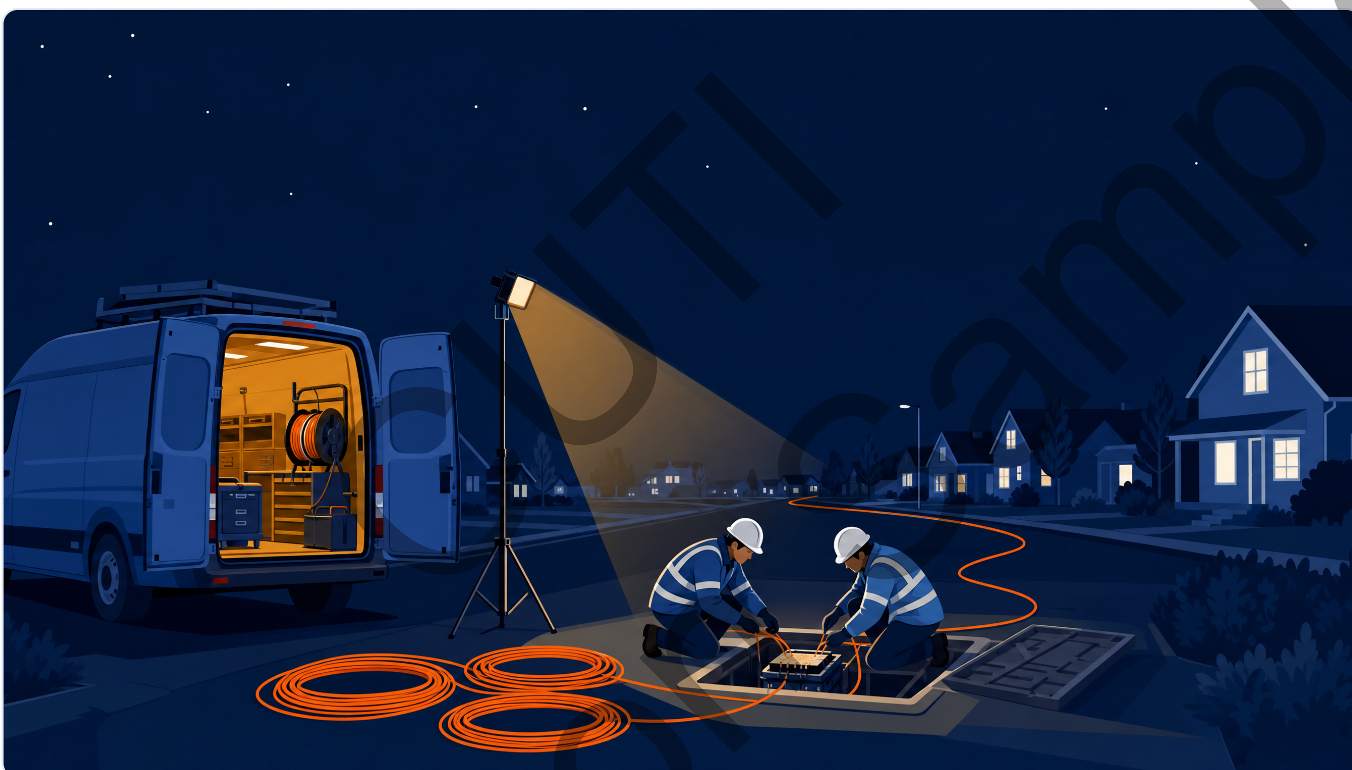

Illustration. The night the design becomes a network.

Learning objectives

After this module, you will know how to:

1. Describe the cable preparation and sealing steps that protect a fusion splice in a splice closure, and identify the defects that cause failures after acceptance.
2. Explain how core-alignment fusion splicing makes a joint, and explain why the estimated loss on the splicer screen is not a measurement.
3. Calculate the bidirectional average splice loss from optical time-domain reflectometer (OTDR) readings, and decide the result against the splice gate.
4. Validate a link with an optical power meter against the M1.1 design value, and decide the result against the link gate.
5. Read an OTDR trace and event table to isolate a fault, and select the test method for a live XGSPON link.

Why this matters

M1.1 gave every fusion splice a budget line of 0.05 dB. M1.3 then put 8 of those splices into 4 splice closures on a real route. This module is where the design meets a technician, an electric arc, and a test set. The numbers a contractor hands back at acceptance become the permanent record of the network. A manager who reads them correctly rejects a defect for the price of one crew visit.

Core concepts

1. Cable preparation and the splice closure

A **fusion splice** is a permanent joint between two fibers, made with an electric arc. The splice itself takes about one minute. The work around it takes most of a day, and it decides whether the joint survives 25 years in a wet handhole.

The crew opens the cable jacket for 2 m to 3 m and cleans the water-blocking material from the buffer tubes. A **ring cut** is a cut around the jacket that releases the jacket without a cut into the buffer tubes below it. The crew clamps the strength members to the closure, so that cable tension never reaches the splices. Cables with metallic armor or metallic strength members get a bond to the closure and a ground at each splice point.

A **splice closure** is the sealed enclosure that holds and protects the splices. Inside it, each buffer tube runs about 1 m to a **splice tray**. About 1 m of bare fiber then sits in the tray around each splice. That stored fiber lets a crew lift a tray and re-splice a joint years later without a new closure. Outside the closure, the **slack loop** stores 10 m or more of cable. A future crew then works in a van instead of in the handhole.

Most outside plant (OSP) closures carry an IP68 rating under IEC 60529, or a rating against Telcordia GR-771-CORE. A sealed closure gets a pressure test before the crew leaves. The crew pressurizes the closure and puts a soap solution on every seam to find leaks. The closure is the part of the network that fails silently. Water enters, the fiber corrodes at a micro-crack, and the fault arrives two winters later as a service call.

Four zones protect one splice

The splice takes a minute. These four zones decide whether it survives 25 years in a wet handhole.

1. Outside

Slack loop

10 m or more of cable, stored on a rack.

A future crew works in a van, not in the handhole.

2. Cable entry

Ring cut and seal

Jacket stripped 2–3 m. The ring cut spares the buffer tubes below it.

Pressure test every seal before you leave.

3. In the closure

Clamp and bond

Strength members are clamped, so cable tension stops there.

Metallic armor is bonded and grounded.

4. Splice tray

Stored fiber

About 1 m of buffer tube to the tray.

About 1 m of fiber around each splice, for a re-splice years later.

Zone 2 is the one that fails silently. Water enters at a bad seal, the fiber corrodes at a micro-crack, and the fault arrives two winters later as a service call.

Figure 1. What a splice closure holds, and where each protection step applies.

[PHOTO: opened splice closure in a handhole, showing clamped strength members, buffer tubes routed to stacked splice trays, and the seal]

[PHOTO: a ring cut on a loose tube cable jacket, beside a correctly stored slack loop on a handhole rack]

SELF-CHECK 1

A crew clamps the cable jacket to the closure but leaves the strength members loose. Name the failure this causes, and name the year it appears.

2. Core-alignment fusion splicing

Illustration. The arc that makes the joint, and the two tools that decide its quality.

The splice sequence is the same for every splice type: strip, clean, cleave. The crew strips the coating from the fiber, cleans the bare fiber with a wipe, and cuts it with a precision cleaver. A **cleave** is a scribe-and-break cut that leaves a flat end face square to the fiber axis. Cleaver makers state a cleave angle of about 0.5° for a good cut. The quality of the cleave sets the quality of the splice, and the splicer rejects a bad cleave before the arc.

Core-alignment fusion splicing aligns the two cores, not the two claddings. Two methods do this. Profile alignment lights the fibers from two directions 90° apart and finds the cores with a camera and software. Local injection and detection couples light into one fiber at a bend. It measures that light at a bend in the other fiber, and moves the stages for maximum transmission. Both methods then run a prefuse cycle to burn off dust, and an arc that fuses the two ends.

The splicer inspects the joint and puts an **estimated splice loss** on the screen. It also protects the joint with a heat-shrink sleeve or a clamp-on protector. A bubble, a core offset, or a bulge in the joint is a reject, and the crew remakes the splice. A repeated arc improves a dark line or spot, but a crew never repeats the arc more than two times.

Six steps, six defects

The splicer catches steps 3 to 5 on the screen. Steps 1, 2, and 6 leave the site undetected.

1. Strip
Remove the fiber coating

Defect: a nicked fiber
Cracks months later. Not visible.

2. Clean
Wipe the bare fiber

Defect: contamination
Becomes a bubble in the arc.

3. Cleave
Scribe and break, about 0.5°

Defect: an angled end face
The splicer rejects it.

4. Align
Match the two cores

Defect: a core offset
Always a remake.

5. Arc
Fuse the two ends

Defect: a bubble or a bulge
Repeat the arc two times at most.

6. Protect
Heat-shrink the sleeve

Defect: an unseated sleeve
The joint moves in the tray.

The cleave sets the result. A splice is never better than the cut that went into it.

Figure 2. Six steps, and the defect each one produces when it goes wrong.

[PHOTO: core-alignment fusion splicer screen showing the X and Y views of a completed splice with the estimated loss value]

SELF-CHECK 2

Name the two alignment methods a core-alignment splicer uses, and state what each one measures.

3. Three numbers for one splice

One splice produces three different numbers, and a manager who treats them as the same number accepts bad work. The three numbers are the estimate, the one-directional reading, and the bidirectional average.

The **estimated splice loss** comes from the splicer. The splicer looks at the joint, or at the light it coupled through the joint, and predicts a value. It has no view of the fiber beyond a few centimeters. An estimate of 0.01 dB is a statement about the picture on the screen, not a measurement of the link.

The **one-directional reading** comes from an OTDR that tests from one end. The OTDR measures loss from backscattered light. The amount of backscatter depends on the **mode field diameter** (MFD) of the fiber, which is the width of the region the light travels in. IEC 60793-2-50 permits a median MFD of 8.8 µm to 9.3 µm for single-mode fiber, so two compliant fibers still differ. When the OTDR passes from a wide-MFD fiber into a narrow-MFD fiber, backscatter rises and the instrument reports less loss than the splice has. In the opposite direction it reports more.

When the backscatter change is larger than the splice loss, the reading goes negative. This event is a **gainer**: an apparent gain of optical power at a splice. No splice adds power. The same splice tested the other way is a **loser**, and reports the true loss plus the full backscatter difference. G.652.D fiber spliced to G.657 bend-insensitive fiber produces this effect often, because the two glass types differ in MFD by design.

The **bidirectional average** is the true splice loss. The crew tests the same splice from both ends and averages the two readings with their signs. The backscatter error is equal and opposite in the two directions, so the average removes it. ANSI/TIA/EIA-455-8 states that OTDR splice loss measurements are taken from both directions and averaged. This is the only method that gives an accurate loss for a splice inside a link.

Number	Source	What it proves	Use in acceptance
Estimated splice loss	The fusion splicer	The joint looks correct and the arc completed.	Crew self-check only. Never an acceptance value.
One-directional reading	OTDR from one end	An event exists at a distance.	Fault location. Never an acceptance value.
Bidirectional average	OTDR from both ends	The true loss of the splice, in dB.	The acceptance value.

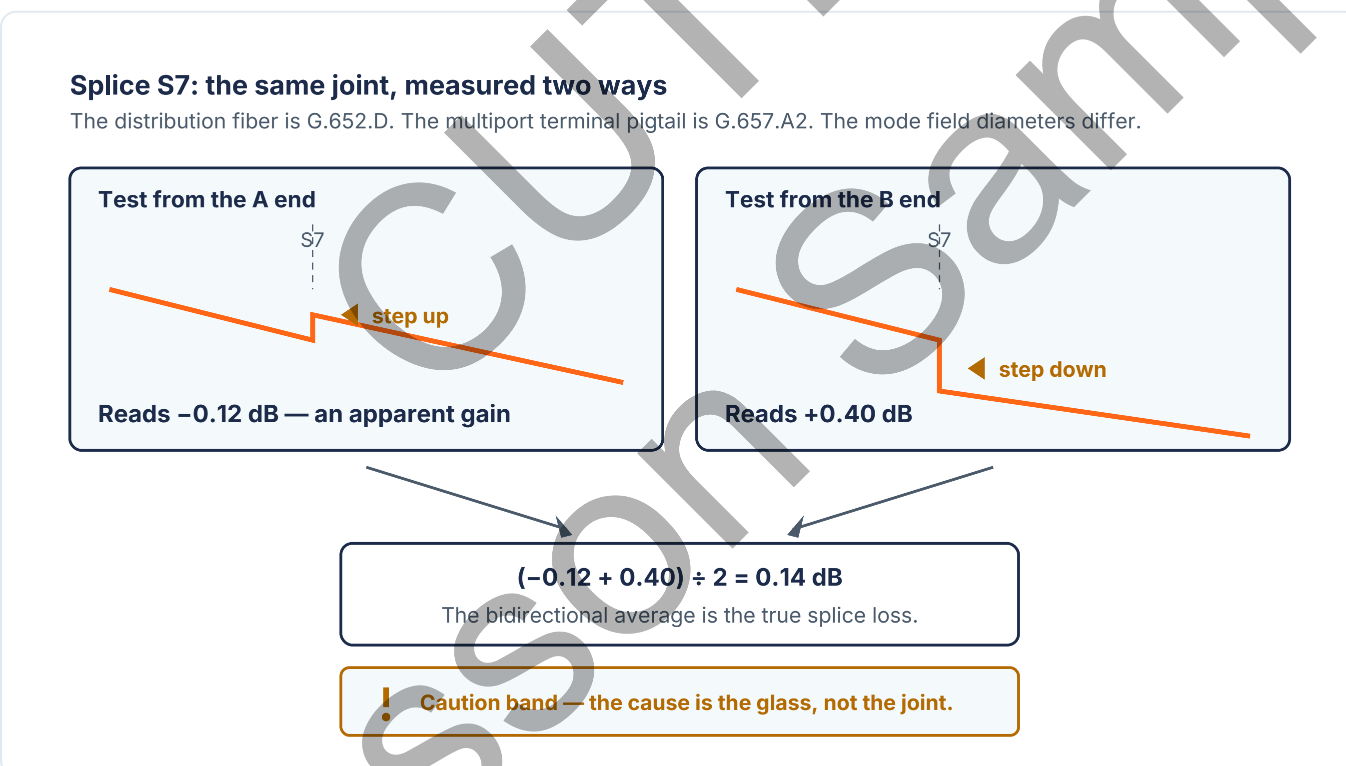

Figure 3. One splice, two readings, one true value. The average removes the backscatter error.

SELF-CHECK 3

An OTDR reads −0.09 dB at a splice from the A end and +0.31 dB from the B end. Calculate the true splice loss, and state why the two readings differ.

4. Acceptance thresholds: the splice gate and the link gate

Two independent gates decide acceptance, and a link passes only when it clears both.

The **splice gate** tests one splice at a time against a threshold. The design target is 0.05 dB, because that is the value M1.1 put in the optical power budget for every splice. The field limits come from published standards. ITU-T L.12 sets 0.2 dB for 90% of single-mode fusion splices and 0.3 dB for all of them. TIA-568 sets a maximum of 0.3 dB for a fusion splice or a mechanical splice. CUTI holds a tighter pass value than either standard, because a residential XGSPON cascade spends most of its budget on the optical splitters.

Bidirectional average	Verdict	Action
0.10 dB or less	Pass	Accept. Record the value in the as-built record.
Above 0.10 dB to 0.30 dB	Caution	Re-splice one time. If the value repeats, record the cause and accept with a note.
Above 0.30 dB	Fail	Re-splice. A value above both ITU-T L.12 and TIA-568 limits never ships.

The caution band exists for one reason: an MFD difference is a property of the two fibers, not a defect in the joint. A second splice at the same joint returns the same value. The crew records the cause, and the manager accepts the value instead of buying four more attempts.

The **link gate** tests the whole link against the design. The optical power budget bought 8 splices at 0.05 dB, so it bought 0.40 dB of splice loss in total. A link with 8 splices at 0.09 dB passes every splice gate. It still spends 0.72 dB, which is 0.32 dB more than the design bought. The link gate catches what the splice gate misses.

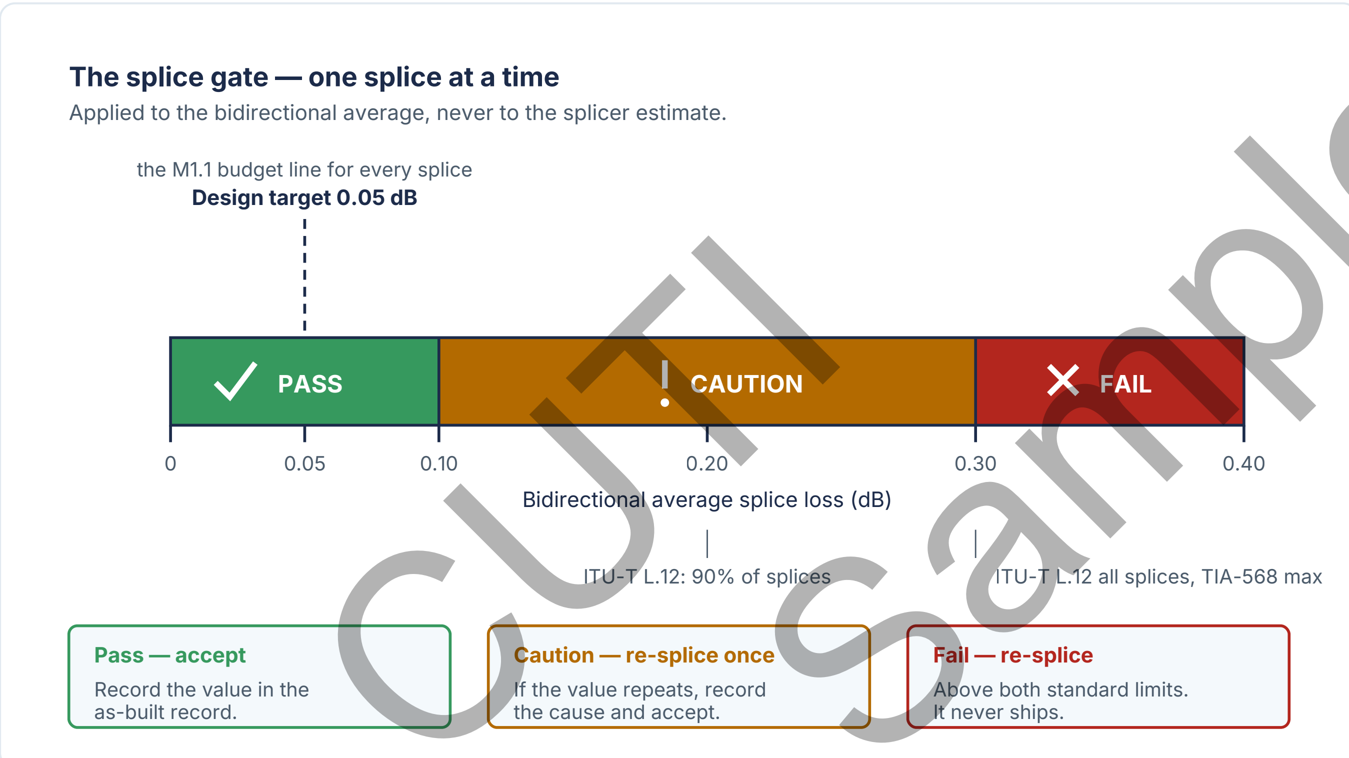

Figure 4. The splice gate. Color and position both carry the verdict.

> **SELF-CHECK 4**
>
> A link has 8 splices. Every bidirectional average is 0.08 dB. State the verdict at the splice gate and the verdict at the link gate, with numbers.

5. Validation with the optical power meter

An **optical power meter** (OPM) measures absolute optical power in dBm. With a matched light source it measures the end-to-end attenuation of a link. This pair is the first tier of validation. It correlates with the equipment, because a source and a meter copy what a transmitter and a receiver do.

Two OPM measurements matter on an XGSPON build. The first is the end-to-end attenuation of the cable plant. A source and a meter measure it at the standard test wavelengths of 1310 nm and 1550 nm. The second is the received optical power (Rx power) at the ONT with the OLT live, measured at the 1577 nm downstream wavelength M1.1 gives. The second measurement is the acceptance value a manager reads.

Note the wavelengths carefully. 1310 nm and 1550 nm are test wavelengths for the fiber. 1270 nm and 1577 nm are the XGSPON service wavelengths from M1.1. A meter set to the wrong wavelength reports a wrong power, because the calibration differs at each wavelength.

The link gate compares the measured value against the design value with a stated tolerance. CUTI accepts a link when the measured Rx power sits within 1.0 dB of the predicted value. On the M1.1 Maple Hill link the predicted value is −19.9 dBm, so the acceptance floor is −20.9 dBm. A reading of −27.5 dBm clears the ONT minimum of −28.0 dBm and fails the link gate by 6.6 dB. That is the M1.1 field scenario, decided by one number.

[PHOTO: optical power meter connected at an ONT with the display showing the received power in dBm at 1577 nm]

SELF-CHECK 5

A link predicts −18.4 dBm. The meter reads −19.1 dBm. Decide the result against the link gate, and give the margin.

6. Validation with the OTDR

An **optical time-domain reflectometer** (OTDR) sends a pulse into the fiber and measures the light that returns. Scattered light comes back along the whole length, and reflected light comes back from connectors and from broken ends. The instrument turns the return against time into a picture of loss against distance. This picture is a **trace**, and each feature in it is an **event**.

A fusion splice is a non-reflective event: a step down with no peak. A connector is a reflective event: a step down with a peak. The slope between events is the attenuation of the fiber in dB/km. The end of the fiber is a large reflection or a fall into noise. The second tier of validation reads this picture, and it is the only tool that assigns a loss to one splice inside a spliced link.

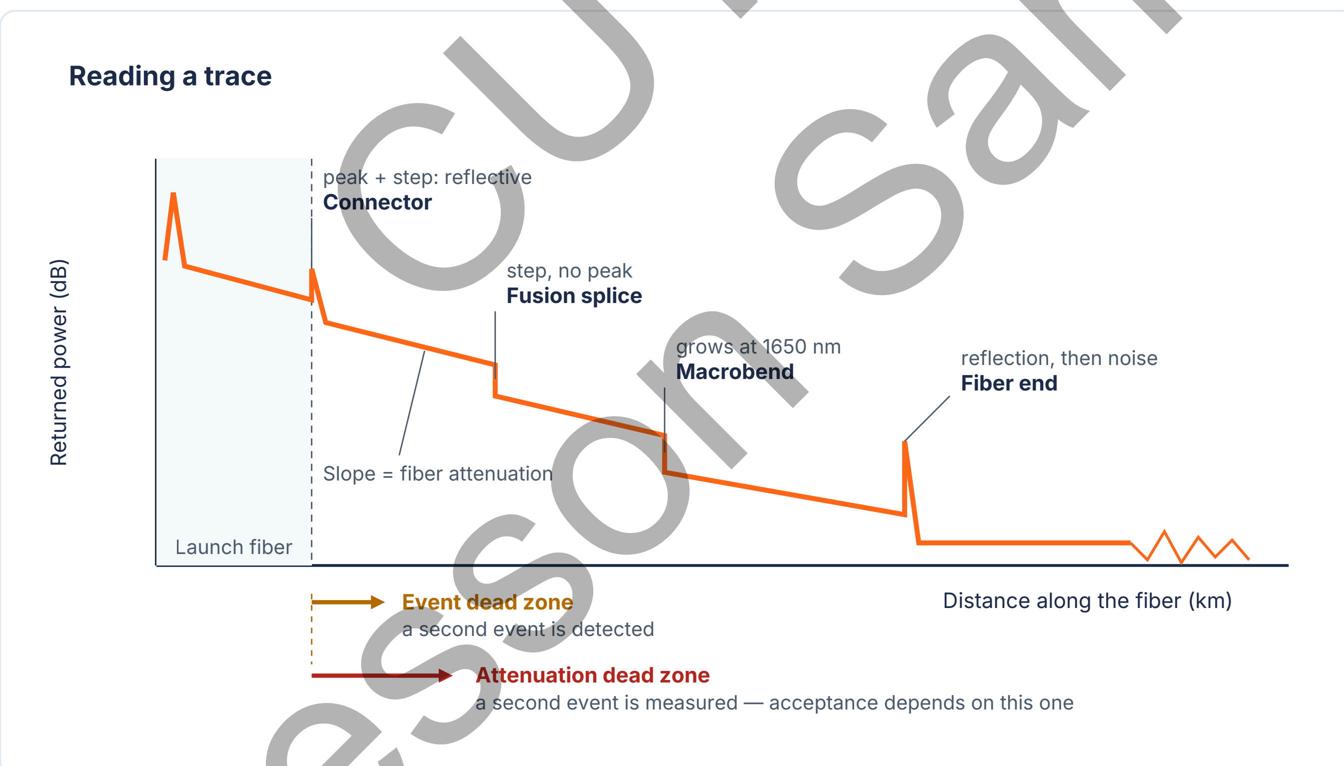

Figure 5. What each feature of a trace means, and where the instrument goes blind.

Where the OTDR goes blind

The test pulse overloads the receiver, so the instrument sees nothing for a distance after every reflective event. Two distances describe this. The **event dead zone** is the distance after a reflective event within which a second event is still detected. The **attenuation dead zone** is the longer distance within which a second event is measured accurately. Acceptance testing depends on the attenuation dead zone, because acceptance needs a loss value and not only a mark on a trace.

A launch fiber solves the problem at the near end. The crew puts a known length of fiber between the instrument and the link. The dead zone then falls inside the launch fiber, and the first connector of the link gets a real measurement. A receive fiber does the same at the far end. An acceptance package with no launch fiber has no measurement of its first connector.

Two more traps sit in every trace. A **ghost** is a reflection that bounced twice and appears at a multiple of a real distance, with an event but no loss. A pinched fiber under a splice tray cover breaks close to the splice. The OTDR then reads the break as one bad splice, and a visual fault locator finds the true fault.

The splitter problem

An XGSPON cascade breaks the OTDR. A test from the OLT passes through the 1×4 and 1×8 optical splitters. The instrument then adds the returns from all 32 branches into one trace. An event appears, and no method assigns it to a branch. Only the end of the longest branch is unambiguous.

Three methods answer this, and a manager selects one before construction starts.

1. **Test each segment before the splitters go in.** This is the cheapest method and it produces the acceptance record for the splices. The feeder and each distribution route get an OTDR trace end to end while the path is still simple.
2. **Test upstream from the subscriber.** After the cascade is complete, an OTDR at the ONT end sees one path and one splitter. The dynamic range needed is high, because

the trace carries the full cascade loss. More averages raise the range without the resolution cost of a wider pulse.

3. **Design test points into the plant.** Connectors on both sides of each optical splitter let a crew test the segments separately for the life of the network. The cost lands in the M1.3 design, not in the test plan.

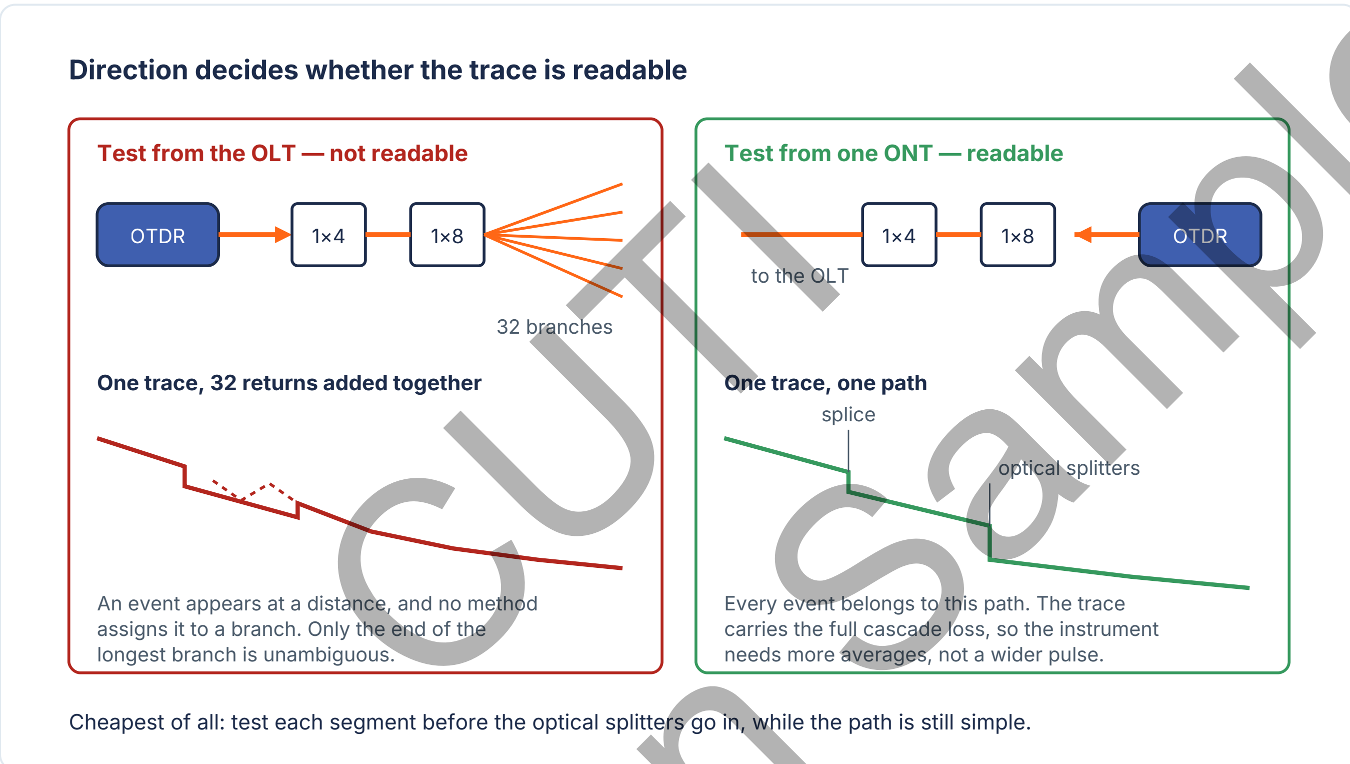

Figure 6. Direction decides whether a trace is readable on a cascade.

A live link adds one more rule. An OTDR pulse at a service wavelength interferes with traffic and reaches the OLT receiver. ITU-T L.66 states that the maintenance test wavelength is a wavelength that carries no communication signal, and it uses 1650 nm in the U-band. A test on a live XGSPON link therefore runs at 1625 nm or 1650 nm through a filter that blocks the service wavelengths. A longer wavelength also responds more strongly to bends, so an out-of-band trace finds a macrobend that a 1310 nm trace hides.

SELF-CHECK 6

A contractor proposes one OTDR test from the OLT after the whole cascade is complete. State two reasons to refuse, and name the test the contractor performs instead.

7. The validation plan and the records that survive the project

A manager writes the validation plan before the first splice, because a test specified after the work is a negotiation. The plan states the gates, the methods, the wavelengths, and the deliverables. Six items make a plan complete.

Item	What the plan states	Why it belongs in the contract
Gates	The splice gate values and the link gate tolerance.	Two gates prevent acceptance of a link that passes on average and fails in one place.
Method	Bidirectional OTDR for splices; source and meter for end-to-end attenuation.	A one-directional package is not a measurement, and it costs nothing to reject it on paper.
Launch and receive fibers	Required on every OTDR test, with lengths stated.	Without them the end connectors carry no measured value.
Wavelengths	1310 nm and 1550 nm for the cable plant; 1577 nm for Rx power; 1625 nm or 1650 nm on a live link.	A wrong wavelength gives a wrong number and a false pass.
Sequence	Segment tests before the optical splitters; end-to-end after.	The splitters make a downstream trace unreadable.

Item	What the plan states	Why it belongs in the contract
Records	Native trace files, the event table, the closure and tray position of every splice, and photographs before closure.	The trace taken today is the reference for every fault for 25 years.

The last row carries the most value and gets the least attention. A trace stored as an image is a picture. A trace stored in the native format of the instrument is data that a future engineer re-analyzes with new markers. The record also states where each splice sits: closure, tray, and position. Without that, a technician in year 7 opens three closures to find one splice. M3.3 takes these records into the as-built process.

> **SELF-CHECK 7**
>
> Name the two records that make a future fault cheap to find, and state what each one gives the technician.

Worked example: accepting the Maple Hill MPT-7 route

The link is the M1.1 link. It runs 6.0 km from the central office (CO) on a 1×32 cascade, with 8 fusion splices in the 4 splice closures M1.3 placed. M1.1 predicts a total attenuation of 21.9 dB and an Rx power of −19.9 dBm, against a usable budget of 27.0 dB. Boreal Splice Services completes the splicing and submits an acceptance package.

Part A — the link gate

The crew measures the Rx power at an ONT on MPT-7 with the OLT live at 1577 nm.

Step	Item	Value
1	Predicted Rx power (M1.1)	−19.9 dBm
2	Acceptance tolerance	1.0 dB
3	Acceptance floor: −19.9 − 1.0	−20.9 dBm
4	Measured Rx power	−20.3 dBm
5	Measured attenuation: 2.0 − (−20.3)	22.3 dB
6	**Link gate: −20.3 ≥ −20.9**	**Pass, 0.6 dB inside the floor**

The measured attenuation of 22.3 dB is 0.4 dB above the design total of 21.9 dB. The link works, and the headroom is 27.0 − 22.3 = 4.7 dB against the 5.1 dB the design bought. A manager who stops here accepts the route.

Part B — the splice gate

The OTDR event table gives both directions for all 8 splices. The average of the two signed readings is the true splice loss.

Splice	Location	A→B (dB)	B→A (dB)	Average (dB)	Verdict
S1	Closure C1, handhole at 1.31 km	0.04	0.06	0.05	Pass
S2	Closure C1, handhole at 1.31 km	0.03	0.05	0.04	Pass

Splice	Location	A→B (dB)	B→A (dB)	Average (dB)	Verdict
S3	Closure C2, at the FDH	0.06	0.04	0.05	Pass
S4	Closure C2, at the FDH	0.05	0.07	0.06	Pass
S5	Closure C3, distribution route	0.04	0.06	0.05	Pass
S6	Closure C3, distribution route	0.41	0.35	0.38	Fail
S7	Closure C4, at MPT-7	−0.12	0.40	0.14	Caution
S8	Closure C4, at MPT-7	0.02	0.04	0.03	Pass
—	**Sum of the 8 averages**	**0.53**	**1.07**	**0.80**	**0.40 dB over budget**

Read the two failures separately. **S6 fails the splice gate.** The average of 0.38 dB is above the 0.30 dB limit in ITU-T L.12 and TIA-568. Both directions agree, so this is a defective joint and not a measurement artifact. **S7 lands in the caution band.** The two readings disagree by 0.52 dB and one is negative, which is the signature of an MFD difference. S7 joins the G.652.D distribution fiber to the G.657.A2 pigtail of the multiport terminal (MPT), so the difference is in the glass. A second splice returns the same value.

The sum of the averages is 0.80 dB against the 0.40 dB the budget bought. The 0.4 dB gap in Part A is exactly this excess. The link gate passed, and the splice gate failed.

Part C — after the repair

The crew re-splices S6 only. The new bidirectional average is 0.08 dB, from readings of 0.07 dB and 0.09 dB. The crew records S7 with its cause and does not touch it.

Step	Item	Value
1	Splice loss before repair	0.80 dB
2	Removed at S6: 0.38 − 0.08	0.30 dB
3	Splice loss after repair: 0.80 − 0.30	0.50 dB
4	Budget line for 8 splices: 8 × 0.05	0.40 dB
5	Residual variance: 0.50 − 0.40	0.10 dB
6	Measured Rx power after repair: −20.3 + 0.30	−20.0 dBm
7	Measured attenuation: 2.0 – (−20.0)	22.0 dB
8	**Headroom: 27.0 − 22.0**	**5.0 dB**
9	**Link gate: −20.0 ≥ −20.9**	**Pass**
10	**Splice gate: 7 pass, 1 caution recorded**	**Pass**

Both gates pass. The route is accepted with a residual variance of 0.10 dB against the design, and that number goes into the as-built record beside S7. The accepted Rx power of −20.0 dBm becomes the reference for every future measurement on this route.

That reference earns its cost two weeks later. In the M1.1 field scenario a slack loop in the handhole at 1.31 km is pressed flat against the wall. The OTDR shows a 7.4 dB step at that distance. The link then measures 22.0 + 7.4 = 29.4 dB, and the Rx power falls to

2.0 − 29.4 = −27.4 dBm. That is the reading the crew found at MPT-7. Acceptance proves a link on the day of the test. Damage after acceptance needs the second gate at turn-up, and it needs the accepted value to compare against.

Field scenario: the acceptance package that tested one way

Illustration. The review that costs one morning and saves eight service calls.

It is two weeks before the Maple Hill turn-up. Boreal Splice Services submits the acceptance package for the MPT-7 route: 8 splices, 8 event readings, every one marked pass. Éric Lavoie, the splicing supervisor, notes that the crew hit 0.05 dB or better on the splicer screen at every joint. S7 recorded a gain, so he entered 0.00 dB.

Dana Okafor, the deployment manager, reads the package and asks four questions. From which end did the crew test? Where is the second direction? Where is the launch fiber on the first connector? What is the sum of the splice losses against the 0.40 dB the budget bought?

The package has one direction, A → B, and no launch fiber. The crew re-tests from both ends the next morning. S6 reads 0.41 dB and 0.35 dB, an average of 0.38 dB. S7 reads −0.12 dB and 0.40 dB, an average of 0.14 dB. The sum of the eight averages is 0.80 dB.

Marc Tessier, the crew lead from NorthLine Construction, proposes re-splicing S6 and S7 together while the closures are open. Éric disagrees on S7. He explains that the distribution cable is G.652.D and the MPT pigtail is G.657.A2, and that the two glass types differ in mode field diameter. A new splice at S7 returns the same 0.14 dB.

[PHOTO: OTDR screen showing an event table with bidirectional readings and one negative value, beside the printed acceptance package]

Decision point. 1. Éric's crew met 0.05 dB on every splicer screen and the link passed the link gate at −20.3 dBm. State why the package is still a reject, in one sentence a subcontractor accepts. 2. Decide which splices get re-spliced, and give the reason for each decision. 3. Name the two clauses you add to the next contract so that this package never arrives again, and state what each clause costs the contractor.

Knowledge check

The question bank for this module is `../assessment/W3-M3.2_bank_v1.md`. It has 15 questions: 3 for each learning objective. Instructors draw quiz forms from the bank. Self-paced students: answer all 15 questions, then attend the review session.

Summary table

Concept	Key point	Number to remember
Splice closure	The closure protects the splice for 25 years; a failed seal is a silent fault.	IP68
Fusion splice	Strip, clean, cleave, align, arc, protect. The cleave sets the result.	Cleave angle ≈ 0.5°
Estimated loss	The splicer screen predicts; it does not measure the link.	Never an acceptance value
Bidirectional average	Average the two signed readings to remove the backscatter error.	(A→B + B→A) ÷ 2
Splice gate	Pass, caution, and fail bands apply to one splice at a time.	0.10 / 0.30 dB
Link gate	Measured Rx power against the M1.1 predicted value, with a tolerance.	1.0 dB

Concept	Key point	Number to remember
OTDR on a cascade	Test segments before the optical splitters, or upstream from the ONT.	32 branches in one trace
Live-link testing	Test out of band so the pulse never reaches the OLT receiver.	1625 or 1650 nm

Key terms

Term	Acronym	Definition
fusion splice	—	A permanent joint between two fibers, made with an electric arc.
splice loss	—	The attenuation one splice adds, in dB.
splice closure	—	The sealed enclosure that holds and protects splices in the outside plant.
splice tray	—	The tray inside a closure that holds the splices and their stored fiber.
ring cut	—	A cut around the cable jacket that releases the jacket without damage to the buffer tubes.
buffer tube	—	The tube inside a loose tube cable that holds a group of fibers.

Term	Acronym	Definition
cleave	—	A scribe-and-break cut that leaves a flat fiber end face square to the axis.
core alignment	—	Splicer alignment of the two fiber cores, not the two claddings.
estimated splice loss	—	The value the fusion splicer predicts after the arc. Not a measurement.
mode field diameter	MFD	The width of the region in which light travels in a single-mode fiber.
gainer	—	An OTDR event that reads as a gain, caused by a backscatter difference.
bidirectional average	—	The average of two signed OTDR readings of one splice. The true splice loss.
splice gate	—	The acceptance test applied to one splice against a threshold.
link gate	—	The acceptance test applied to the whole link against the design value.
acceptance value	—	The design value a measurement is compared against at turn-up.
trace	—	The OTDR picture of loss against distance along a fiber.
event	—	A feature in a trace produced by a splice, a connector, a bend, or an end.

Term	Acronym	Definition
reflective event	—	An event with a reflection peak, typical of a connector.
event dead zone	—	The distance after a reflective event within which a second event is detected.
attenuation dead zone	—	The distance after a reflective event within which a second event is measured accurately.
launch fiber	—	A known length of fiber before the link that moves the dead zone off the first connector.
receive fiber	—	A known length of fiber after the link that gives the last connector a measured value.
ghost	—	A false event from a double reflection, with no loss at the event.
out-of-band testing	—	Testing at a wavelength that carries no service traffic, for work on a live link.
visual fault locator	VFL	A visible light source that finds breaks and bends close to a connector or splice.
optical power meter	OPM	A test set that measures absolute optical power in dBm.
optical time-domain reflectometer	OTDR	A test set that locates and measures attenuation events along a fiber.

Self-check answers

Self-check 1. Cable tension reaches the splices instead of the closure. The fibers break or the splices pull apart at the first thermal cycle. The fault appears in the first winter after construction.

Self-check 2. Profile alignment lights the fibers from two directions 90° apart and finds the core positions with a camera and software. Local injection and detection couples light in at a bend and measures the transmitted light at a bend in the other fiber.

Self-check 3. The true splice loss is $(-0.09 + 0.31) \div 2 = 0.11$ dB. The two fibers have different mode field diameters, so backscatter rises in one direction and falls in the other. The error is equal and opposite, and the average removes it.

Self-check 4. Every splice passes the splice gate, because 0.08 dB is below the 0.10 dB pass value. The link fails the link gate on splice loss. The 8 splices spend $8 \times 0.08 = 0.64$ dB against the 0.40 dB the budget bought, which is 0.24 dB of extra attenuation.

Self-check 5. Accept. The difference is 0.7 dB, which is inside the 1.0 dB tolerance, and the margin to the acceptance floor of −19.4 dBm is 0.3 dB.

Self-check 6. First, the two optical splitters combine the returns of 32 branches, so no event is assigned to a branch. Second, a pulse at a service wavelength on a live link interferes with traffic and reaches the OLT receiver. The contractor tests each segment before the optical splitters go in, and tests upstream from an ONT afterward.

Self-check 7. The native trace file gives the technician a reference to compare a new trace against, with the original markers and the option to re-analyze. The splice position record gives closure, tray, and position, so the technician opens one closure instead of three.