HFCL

Whitepaper

Hollow-Core Fibre: Nested Anti-Resonant Nodeless Fibre

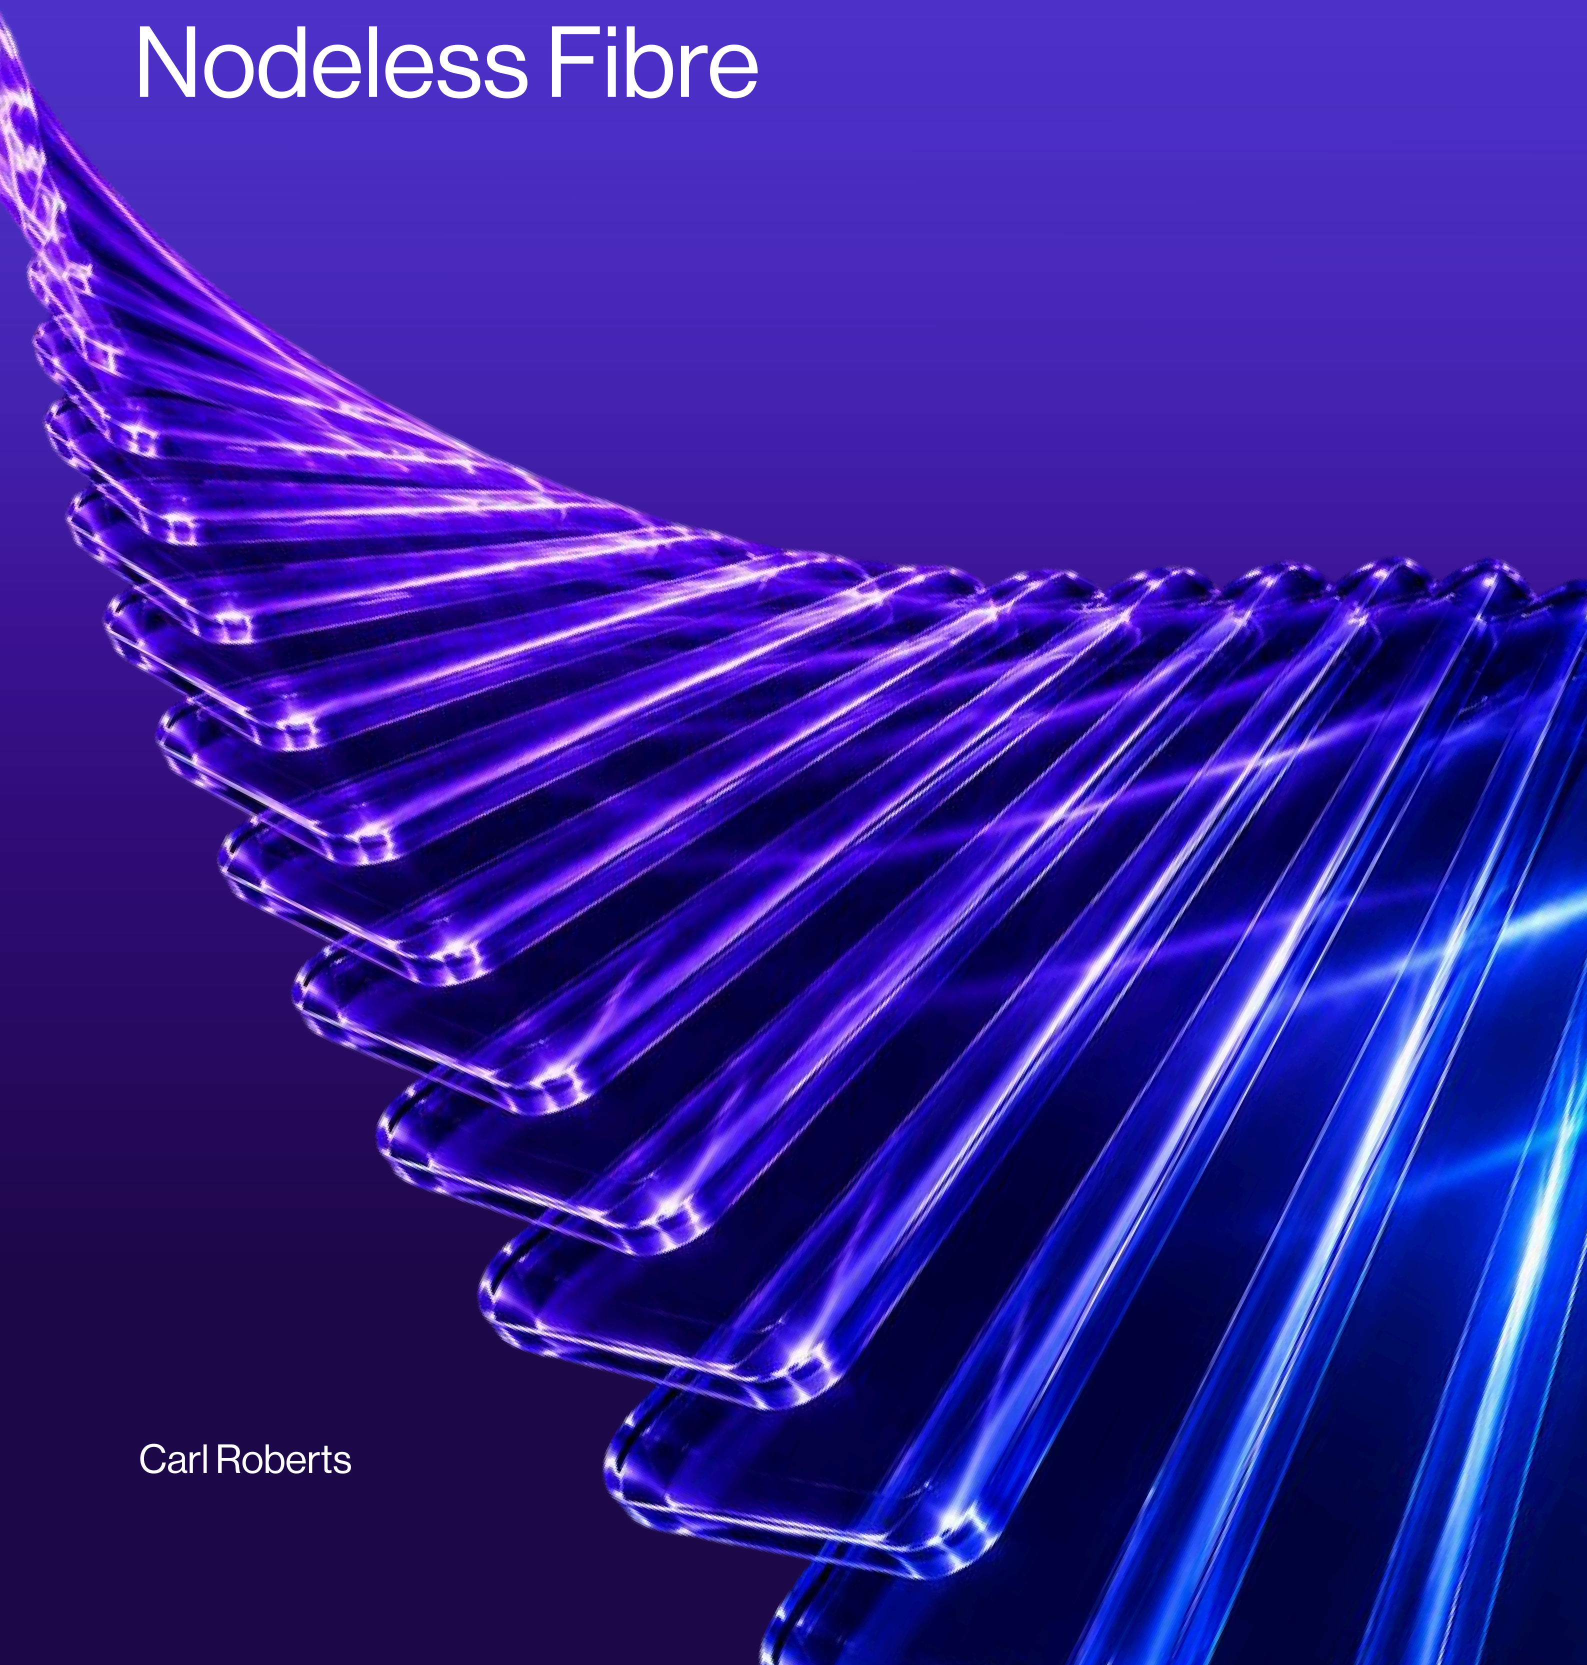

Carl Roberts

HFCL

Contents

HFCL

Introduction

The global telecommunications industry has spent more than four decades refining a single guiding principle: that light should be transmitted through a strand of solid silica glass, confined by the difference in refractive index between a denser core and a less-dense surrounding cladding. This approach of total internal reflection has been extraordinarily successful, enabling internet speeds and capacities that would have seemed impossible thirty years ago. Yet it carries a fundamental physical limitation: the light must travel through glass, and glass imposes irreducible penalties on every signal that passes through it.

These penalties include Rayleigh scattering (where microscopic density variations in the glass deflect photons out of the fibre), chromatic dispersion (where different wavelengths travel at slightly different speeds, blurring pulses over distance), nonlinear optical effects (where photons at high power interact with the glass itself, generating noise and crosstalk), and a fundamental speed ceiling where light in glass travels at only ~68% of its vacuum speed, adding unavoidable latency to every kilometre of fibre. Together, these intrinsic glass limitations bear directly on two markets.

The first is ultra-low-latency financial trading, where sub-millisecond advantages translate into competitive revenue. The second is AI infrastructure, where propagation delay across campus- and metro-scale interconnects contributes to GPU synchronisation overhead.

Hollow core fibre removes all of these penalties at a stroke. By guiding light through an air-filled hollow core rather than through solid glass, HCF simultaneously breaks the Rayleigh scattering floor, suppresses nonlinear effects by four orders of magnitude, reduces latency by approximately 30%, and opens transmission windows in wavelength bands where silica is simply opaque. The enabling mechanism is not total internal reflection but a completely different optical phenomenon: anti-resonant reflection.

This white paper focuses specifically on Nested Anti-Resonant Nodeless Fibre (NANF), the fibre architecture that has achieved the lowest losses ever recorded in any optical waveguide, and the design that underpins the majority of current commercial HCF deployments. It explains how NANF works from first principles and covers the physics of anti-resonant reflection.

2 How NANF Works

2.1 The Core Idea: Guiding Light Through Air

In a conventional single-mode fibre (SMF), the glass core has a slightly higher refractive index than the surrounding cladding. Light travelling down the fibre at a sufficiently shallow angle undergoes total internal reflection at the core–cladding boundary and is unable to escape. It bounces down the fibre, staying confined within the solid glass core.

NANF takes a fundamentally different approach. The core is hollow and contains nothing but air (or in some variants, an inert gas). Light has no denser medium to reflect from in the conventional sense. Instead, it is confined by a ring of thin-walled glass capillary tubes positioned around the hollow core.

These tubes do not rely on standard reflective coatings or index changes; instead, they use phase interference to transform each thin glass wall into an ultra-low-loss reflection boundary, bouncing light back into the core through the interference mechanism described in Section 3.

Because the light travels almost entirely through air, it inherits the properties of air rather than glass. Air has negligible optical absorption, near-zero nonlinearity, and a refractive index of essentially 1.000, meaning light travels through NANF at 99.7% of its vacuum speed, compared with only ~68% in conventional SMF. This 1.46× speed difference translates directly into a ~30% reduction in signal propagation latency over any given distance.

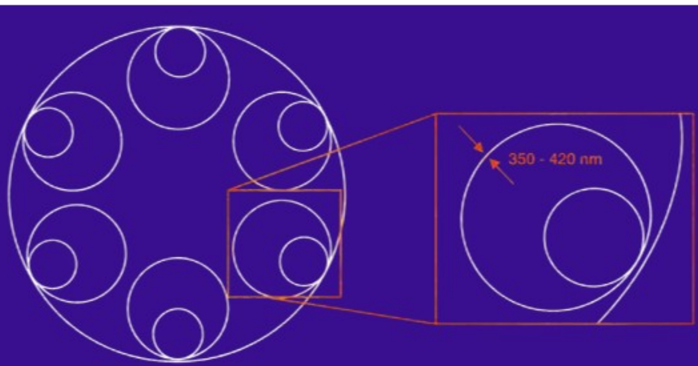

Figure 1: NANF Cross-Section: six nodeless capillary tubes arranged around a hollow air core. The wall of each capillary is 350–420 nm thick for C-band operation, tuned to reflect the operating wavelength back into the core via anti-resonant interference. The tubes do not touch each other (nodeless), which eliminates a major source of mode coupling loss compared to earlier anti-resonant designs.

2.2 NANF Architecture

NANF belongs to the anti-resonant HCF family. Its architecture consists of three structural elements:

- The hollow core: A central air-filled region, typically 20 – 50 µm in diameter, through which the signal light travels. The core must be free of any contamination, dust or moisture ingress creates localised scattering and loss spikes.
- The anti-resonant capillary tubes: A ring of typically five or six thin-walled hollow glass tubes positioned immediately around the core. Each tube wall is tuned to a specific thickness selected to create anti-resonant reflection at the operating wavelength.
- The solid glass outer cladding: A solid silica jacket that provides mechanical support, houses the capillary tubes, and carries the polymer coating. The cladding is not involved in the optical

2.3 NANF vs Earlier HCF Designs

Understanding why NANF achieves such low losses requires a brief comparison with the HCF designs that preceded it:

Design	Guiding mechanism	Best loss (dB/km)	Multimode Fiber (MMF)
Photonic Bandgap (PBGF)	Photonic bandgap	~1.2	Narrow bandwidth; surface mode coupling at core boundary
Simple AR tube	Anti-resonance	~0.3	Contact nodes between tubes cause higher-order mode loss
NANF	Anti-resonance (nodeless / double-nested)	0.091	Surface scattering at core–tube interface; now dominant
ST-HCF	Anti-resonance (support tube)	0.04	Lowest reported minimum attenuation; not directly comparable with the 0.091 dB/km average

The key architectural innovation in NANF was removing the contact nodes between adjacent capillary tubes which had a disproportionately large effect on loss. In earlier designs with touching tubes, light could couple into the glass nodes and leak into the cladding. NANF's nodeless geometry eliminates this leakage path, reducing loss by a factor of three compared to simple anti-resonant designs. DNANF (Double-Nested ANF) takes this further by nesting a second, smaller capillary tube inside each outer capillary, further suppressing mode leakage and achieving 0.091dB/km, the lowest average loss reported for any optical fibre. The separate 0.04dB/km figure is a minimum attenuation reported by YOFC for a support-tube architecture, not a DNANF result.

2.4 Single-Mode Operation

A fibre is single-mode when only one spatial mode of light can propagate stably down its length. In HCF, the hollow core naturally supports multiple modes, but higher-order modes experience significantly higher loss than the fundamental mode because they interact more strongly with the capillary tube walls. Over any significant propagation distance, the higher-order modes are filtered out, leaving only the fundamental mode. This effectively single-mode operation is sufficient for telecommunications and data centre applications and is one of NANF's strong practical advantages over other HCF designs.

3. How Anti-Resonant Reflection Works

3.1 The Principle

The term anti-resonance refers to a condition in which two reflected light waves combine to produce a stronger-than-normal reflection. To understand this, it helps to first understand thin-film interference, the same phenomenon that produces the iridescent colours in soap bubbles and oil films on water.

When a beam of light strikes a thin sheet of glass, it does not simply pass straight through or reflect from a single surface. Two separate reflections occur simultaneously: one from the front surface of the glass (where the light first enters), and one from the back surface (the far side of the sheet). These two reflected beams travel back in the same direction, but they have travelled different distances; the second beam has crossed the glass thickness twice more than the first. This path difference means that the peaks and troughs of the two light waves may or may not line up when they recombine.

Whether the waves line up depends on the relationship between the glass thickness and the light's wavelength. Two scenarios are possible:

- Anti-resonant (constructive interference): The path difference is such that the two reflected waves arrive back in step, peak aligns with peak, trough with trough. They reinforce each other, producing a combined reflected beam that is much stronger than either reflection alone. Nearly all the incoming light bounces back into the air core. This is the desired operating condition for a NANF capillary wall.

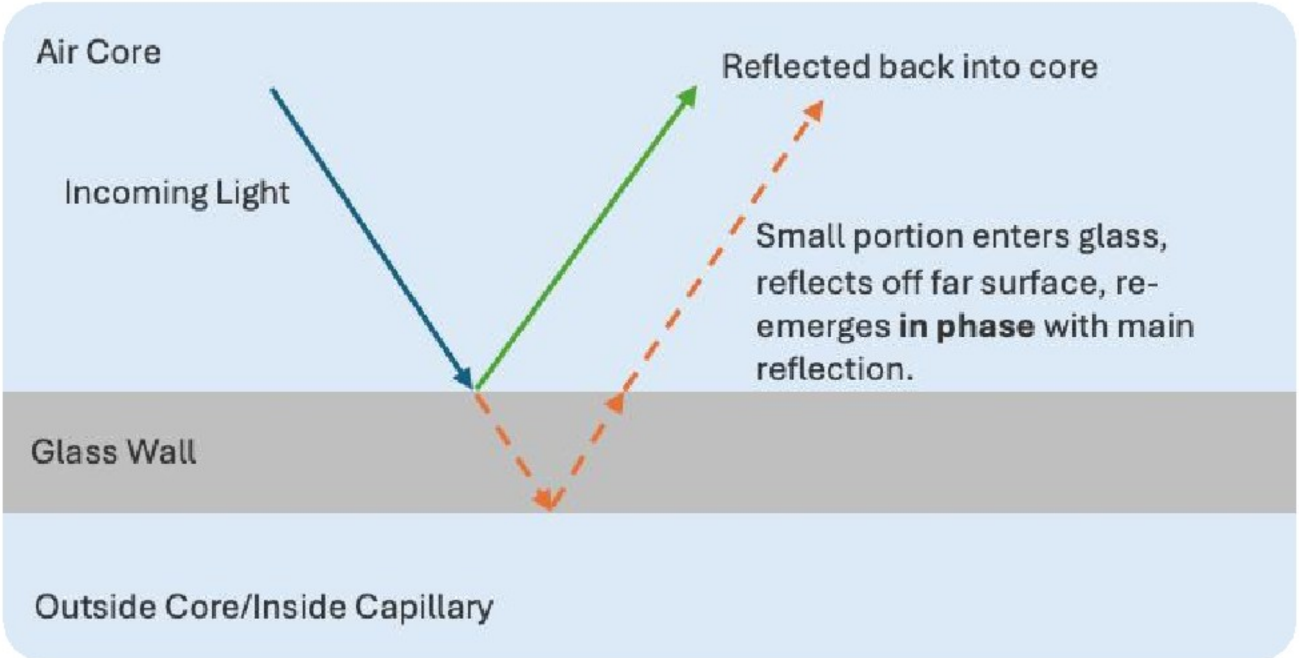

Figure 2: Anti-resonant reflection at a NANF capillary wall. An incoming ray from the core strikes the glass wall. A large fraction reflects from the front surface. A small fraction enters the glass, reflects off the far surface, and re-emerges parallel to the first reflection. When the wall thickness is tuned so that these two beams are in phase, they reinforce each other (constructive interference), creating a near-perfect mirror that sends almost all the light back into the core.

- Resonant (destructive interference): The path difference causes the two reflected waves to arrive exactly out of step; the peak of one aligns with the trough of the other. They cancel each other, and almost no light is reflected back. Instead, the light passes through the glass wall and leaks out of the core. This is the condition the NANF designer works to avoid at the operating wavelength.

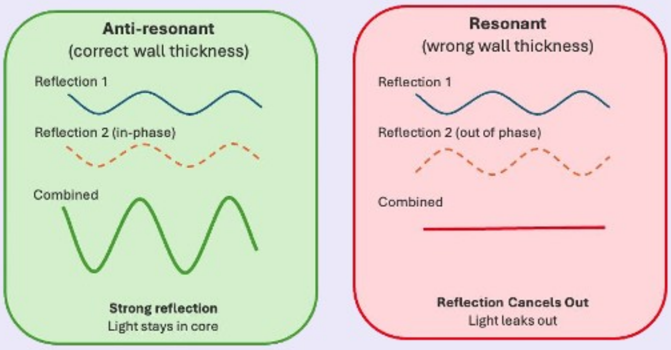

Figure 3: The two possible interference outcomes. Left: anti-resonant condition where - the two reflected waves are in phase and reinforce, producing strong reflection and effective light confinement. Right: resonant condition - the two reflected waves are out of phase and cancel, destroying the reflection and causing light to leak through the wall. The NANF wall thickness is chosen so that the operating wavelength always falls in the anti-resonant regime.

3.2 Mathematical Conditions

The anti-resonant condition is satisfied when the capillary wall thickness t meets:

t = (2m – 1) · λ / (4 · √(n2 – 1))

m: anti-resonance order (integer: 1, 2, 3 ...); m = 1 gives the thinnest wall
λ: operating wavelength (e.g. 1550 nm for the telecom C-band)
n: refractive index of the glass wall (n ≈ 1.45 for silica)
√(n² – 1): the effective index term for a wall struck at near-grazing incidence, as it is from an air core; for silica this is ≈ 1.05, not n ≈ 1.45

The √(n² – 1) term is what distinguishes a hollow-core wall from an ordinary thin film. Light in the air core travels almost parallel to the wall and strikes it at near-grazing incidence, not near-normal incidence, so the relevant quantity is not the refractive index itself but the transverse component of the wave-vector inside the glass. For silica this evaluates to 1.05 rather than 1.45, a difference of nearly 40% in the resulting wall thickness. For a NANF designed for the telecom C-band (λ = 1550 nm) in silica, the first-order (m = 1) anti-resonant wall thickness is therefore:

t ≈ 1550 / (4 × 1.05) ≈ 369 nm

Higher orders satisfy the same condition at the same operating wavelength: m = 2 requires t ≈ 1107 nm and m = 3 requires t ≈ 1845 nm. The choice between them is a trade between manufacturability and spectral width. A thicker wall is more forgiving to draw and mechanically more robust, but it packs the resonance comb more tightly. At m = 1 the resonances fall at 775 nm and below, leaving an open-ended window above; at m = 2 they fall at 1162 nm and 2325 nm, limiting the usable band into that span. At m = 3 a resonance lands at approximately 1292 nm, inside the O-band. Manufacturers therefore work at m = 1, and commercial C-band NANF core walls sit in the region of 350 – 420 nm. The range reflecting the exact band targeted and the tolerance budget rather than any difference of principle. Section 3.4 examines the resulting spectral windows in further detail.

3.3 Why the Glass Wall Must Be So Thin

The anti-resonance condition links wall thickness directly to operating wavelength. For operation around the C-band a NANF needs roughly 370 nm of glass, about a quarter of a micron, held uniformly around a core 20 to 50 μm across, along every kilometre of fibre drawn.

The thinness is the result of wanting the entire resonance comb pushed below the operating band.
The complementary condition, the thickness at which the wall becomes effectively transparent and the light escapes, is:

$$t = m \cdot \lambda / (2 \cdot \sqrt{(n2 - 1)})$$

For silica at 1550 nm this places the first resonance at t ≈ 738 nm, exactly twice the m = 1 anti-resonant thickness. The wall must therefore sit near the midpoint of that interval; any thickness error shifts the spectrum rather than switching the guidance mechanism off. Rearranging the condition, a wall of thickness t places its resonances at $\lambda m = 2t\sqrt{(n2 - 1)} / m$, so at first order the resonance edge moves roughly 2.1 nm in wavelength for every nanometre of wall thickness. A 370 nm wall puts that edge near 775 nm; a 10 nm error moves it to about 796 nm, which is ~700 nm clear of the C-band and of no consequence in the middle of a window more than a thousand nanometres wide. A wall a few nanometres off target is not, by itself, what ruins a NANF.

What matters is variation, and on three counts. Along the fibre, thickness that drifts or oscillates presents the guided mode with a continuously changing boundary, coupling power into the cladding and into higher-order modes. The penalty accumulates over distance even though no point on the fibre is anywhere near resonance. Across the cross-section, the capillaries must match one another, or the core mode sees an asymmetric mirror and the modal and polarisation purity that NANF is valued for degrades. And between the core-surround capillaries and the nested tubes, the two resonance combs must be held either coincident or deliberately apart, as Section 3.4 describes. Achieving all three at drawing speeds of several hundred metres per minute, continuously over many kilometres, is one of the central manufacturing challenges in NANF production.

3.4 Bandwidth and Spectral Windows

Anti-resonant guidance is intrinsically broadband, and this is one of its principal advantages over photonic bandgap HCF, in which light is confined only within a narrow bandgap typically a few tens of nanometres wide. An anti-resonant wall transmits strongly only at the discrete resonant wavelengths, everywhere between them it reflects. For a wall of thickness t those resonances fall at:

$$t = m \cdot \lambda / (2 \cdot \sqrt{(n2 - 1)})$$

A single wall thickness therefore produces a comb of narrow high-loss lines with broad low-loss windows in the gaps between them. Because successive resonances are spaced as 1/m, those windows are not of equal width: they widen progressively towards longer wavelengths, and the fundamental window above λ_1 is open-ended, bounded on its long-wavelength side by other loss mechanisms discussed below. Designing the spectral response of a NANF is, in the first instance, simply a matter of pushing every resonance out of the band of interest.

Figure 4 shows why the wall has to be thin. Modelled for a wall 1200 nm thick, resonances land at roughly 2520, 1260, 840, 630 and 504 nm. The telecom region is fragmented, with a high-loss spike sitting squarely in the O-band and a second approaching the long-wavelength edge of the L-band. No continuous window of useful width survives. Figure 5 shows the same model for a wall of 350 nm, dimensioned for C-band operation. Only two resonances now fall within the plotted range, at approximately 735 nm and 368 nm, both far below the telecom bands. The result is a single unbroken window running from around 800 nm to beyond 2000 nm, more than 1000 nm of continuous spectrum, set against the 35 nm of the C-band and roughly 90 nm for C+L combined. One NANF design covers the O, E, S, C, L and U bands simultaneously, with a great deal of spectrum still unused.

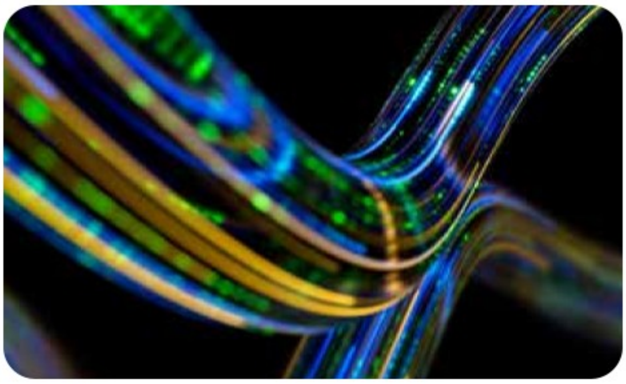

Two consequences follow that matter commercially. First, because the guided field travels in air rather than glass, there is no water absorption peak: the E-band around 1360 - 1460 nm, which legacy G.652 fibre effectively lost and which even low-water-peak G.652.D only partially recovers, is fully available in NANF. This assumes the core is dry and stays dry: water vapour or hydrogen ingress into the hollow core would reintroduce absorption, which is one reason hermetic termination (Section 6.3) matters operationally. Second, the addressable spectrum ceases to be a property of the cable. On standard single-mode fibre the usable band is bounded by the fibre's own attenuation profile as much as by amplifier availability. In NANF, the constraint moves almost entirely to the transceiver and amplifier. As S-band TDFAs and O-band semiconductor amplifiers mature, further capacity can be lit on already-installed NANF without touching the duct, a meaningful argument on congested urban routes and long permitting cycles, where civil works dominate the installed cost far more than the fibre itself.

The anti-resonant window should not however, be confused with the low-loss window, and both figures make the distinction visible: within the broad window the loss is not flat but forms a shallow trough. Two mechanisms set its sides. Towards shorter wavelengths, surface scattering from the frozen-in capillary-wave roughness of the glass - air boundaries rise steeply, scaling as λ^{-3}. Towards longer wavelengths, leakage through the finite capillary structure grows, and the infrared absorption edge of silica begins to couple into the small fraction of the field that overlaps the glass. The minimum sits between the two, typically in the 1.5 - 2.0 µm region for current NANF geometries. This coincides with the bands the industry already operates in. How deep that minimum goes is governed by core diameter, tube count and nesting geometry, as captured by the Zeisberger–Schmidt leakage model; the wall thickness decides only where along the spectrum the windows fall.

One practical qualification applies to real fibre. A NANF has more than one glass wall: the core-surround capillaries and the nested tubes within them each have their own thickness, and so each generates its own resonance comb. If the two combs are not aligned, resonances belonging to the nested tubes intrude into what would otherwise be a clean window. Because the nested tubes carry far less field than the core wall, these appear as shallow features rather than full loss spikes. Manufacturers therefore either hold both wall thicknesses to the same target or deliberately detune the nested tubes so that their resonances fall outside the operating band, which adds a second dimensional tolerance to an already demanding draw.

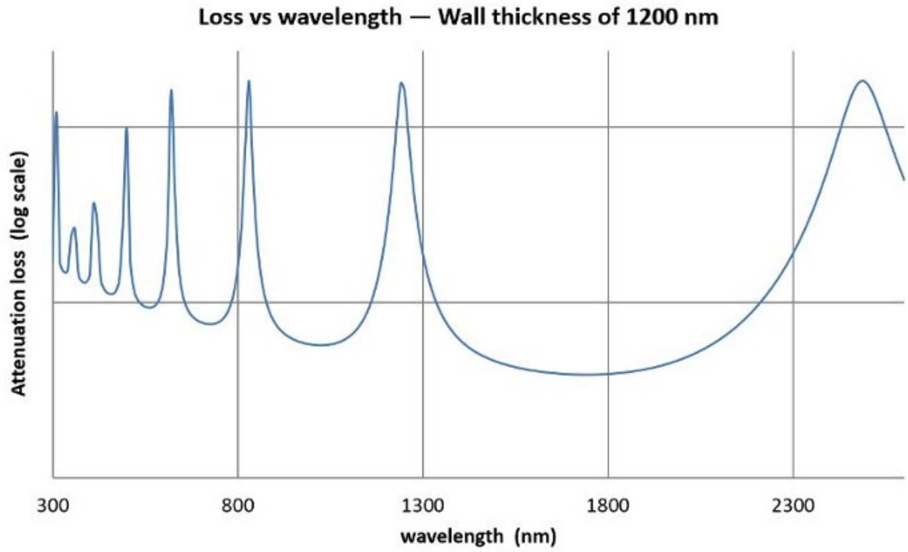

Figure 4 Modelled attenuation against wavelength for an anti-resonant wall 1,200 nm thick (logarithmic loss axis). Resonances fall at approximately 2520 nm, 1260 nm, 840 nm, 630 nm and 504 nm, breaking the telecom region into narrow fragments and placing a high-loss line directly in the O-band. No usable continuous window spans the C and L bands.

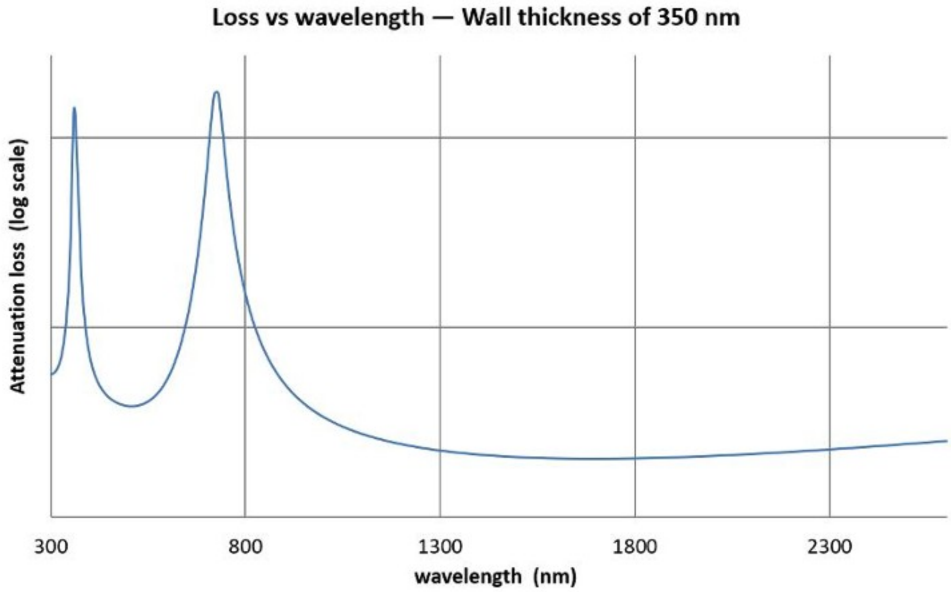

Figure 5 — The same model for a wall 350 nm thick, dimensioned for C-band operation. Only two resonances remain in range, at roughly 735 nm and 368 nm, both well below the telecom bands, leaving a single unbroken window from about 800 nm to beyond 2 µm. The shallow minimum near 1.6 µm is set by the balance between short-wavelength surface scattering and long-wavelength leakage, not by the anti-resonance condition.

4. Attenuation in HCF

4.1 Rayleigh Floor

The fundamental limit on attenuation in conventional SMF is set by Rayleigh scattering, an intrinsic property of glass. Random density fluctuations, frozen into the silica when it solidified during drawing, scatter a small fraction of passing photons in all directions. This loss scales as $1/\lambda^4$, so it is stronger at shorter wavelengths (which is why standard SMF operates at 1310 nm or 1550 nm, long-wavelength windows where Rayleigh scattering is minimised). Even at 1550 nm, the theoretical minimum attenuation for silica SMF is approximately 0.14 – 0.15 dB/km. Every SMF-based system in the world is limited by this floor.

HCF surpasses this limit entirely. Because light travels through air it does not experience Rayleigh scattering from glass density fluctuations. The fraction of light that overlaps with the glass in a well-designed NANF is only around 0.001 – 0.01% (10^{-5} – 10^{-4}) of the total mode energy. Rayleigh scattering from such a tiny glass-overlap fraction contributes a negligible amount to overall attenuation. This is why Hollow Core Fibre (HCF) technology has already achieved a minimum attenuation of 0.04 dB/km announced by YOFC at MWC Barcelona in March 2026 and presented at OFC 2026, less than one-third of the SMF minimum. Researchers argue that no fundamental physical barrier prevents HCF attenuation being pushed towards 0.01 dB/km.

4.2 Attenuation in NANF

While HCF eliminates Rayleigh scattering as the dominant loss mechanism, it introduces new loss mechanisms specific to its structure. These are summarised below:

Loss Mechanism	Typical contribution	Dominant in?	Mitigation
Surface scattering	0.03 - 0.08 dB/km	In NANF	Smoother core–tube boundary; DNANF nesting suppresses mode overlap
Confinement loss	0.001 - 0.02 dB/km	Moderate	More capillary tubes, larger core, thicker outer jacket
Macrobend / bend loss	0.001 - 0.05 dB/km	Installation	Bend radius >20 mm (installation), >10 mm (operation), improved cable design
Residual glass absorption	<0.001 dB/km	Negligible	Ultra-high-purity silica, contributes <1% of total loss
Rayleigh scattering (glass)	~0.0001 dB/km	Negligible	Inherently suppressed, light barely touches glass

Surface scattering is currently the dominant loss mechanism in NANF. The inner surface of the capillary tubes where they border the hollow core is not perfectly smooth at the atomic scale. Light grazes this surface at a very shallow angle as it propagates, and any roughness deflects a fraction of it out of the guided mode. This is fundamentally different from Rayleigh scattering: it is a surface phenomenon rather than a bulk one, and it scales as λ^{-3} rather than λ^{-4}. The key implication is that the loss of NANF is less wavelength-sensitive than SMF. NANF remains low-loss across a broader wavelength range and the path to lower loss lies in improving surface smoothness during drawing rather than changing glass composition.

Confinement loss is the fraction of mode energy that leaks through the capillary tube walls and escapes into the outer cladding. It is controlled primarily by the core size, the number and geometry of tubes, and the number of glass layers surrounding the core. Increasing the core diameter reduces confinement loss but increases bend sensitivity; NANF designers balance these trade-offs numerically. DNANF's nested tube architecture adds an additional layer of confinement, reducing leakage compared with single-tube NANF.

Bend loss is significantly higher in HCF than in SMF. When the fibre is bent, the effective path difference between the inside and outside of the bend distorts the mode and allows more light to leak through the capillary walls. Minimum bend radii of 20 - 30 mm during installation and 10 - 20 mm in operation are typical for current NANF designs. Cable and duct design must account for this, as sharp corners during installation are a significant risk that does not exist to the same degree with SMF.

4.3 Attenuation Progression

The following table documents the key attenuation milestones for NANF and related HCF designs. The pace of improvement from 1.3 dB/km in the first modern AR-HCF demonstrations to 0.04 dB/km in less than a decade is remarkable and mirrors the early trajectory of SMF improvement in the 1970s and 1980s. The difference is that NANF has already surpassed the ultimate physical floor that SMF cannot breach.

Date	Organisation	Record loss	Fibre type	Significance
2018	Univ. Southampton	1.3 dB/km	NANF	Early NANF demonstration
OFC 2022	Univ. Southampton	0.174 dB/km	DNANF	First sub-0.2 dB/km hollow core fibre; first DNANF record
OFC 2024	Univ. Southampton	0.11 dB/km	DNANF	Within striking distance of silica Rayleigh floor (0.14 dB/km)
Sep 2025	Southampton / Microsoft	0.091 dB/km	DNANF	First HCF to break the silica Rayleigh floor: 0.091 dB/km at 1550 nm, remaining below 0.2 dB/km across a 66 THz window (Nature Photonics, 1 Sep 2025)
OFC 2025	YOFC	0.05 dB/km	ST-HCF	World record at time of announcement; 21.7 km drawn length
MWC Barcelona / OFC 2026	YOFC	0.04 dB/km	Advanced AR-HCF	Lowest reported minimum attenuation; vendor-reported, not yet peer-reviewed
Production 2026	Microsoft / Lumenisity	0.091 dB/km (research grade); ~0.10–0.15 dB/km typical production	DNANF	0.091 dB/km is the Nature Photonics lab result; 1,200+ km deployed in Azure; typical production loss is 0.10–0.15 dB/km at scale

Two numbers in this table deserve particular attention. First, the 0.091 dB/km research-grade result from Microsoft/ Southampton published in Nature Photonics demonstrates that sub-silica-floor loss is achievable in a real fibre. The 1,200+ km of live Azure infrastructure carries traffic at this design's specification, though typical production yields across the full deployed base were ~0.10 - 0.15 dB/km. Second, the three headline figures measure different things and should not be read as a single series: 0.040 dB/km is the lowest reported minimum attenuation, 0.091 dB/km the lowest reported average loss at 1550 nm, and 0.10 - 0.15 dB/km the range currently achieved in volume production. It is the gap between the 0.091 dB/km laboratory result and the 0.10 - 0.15 dB/km production range that quantifies the remaining manufacturing challenge, and closing it is likely the single most commercially significant target for the NANF industry over the next five years.

4.4 Comparison with SMF Attenuation

Attenuation quick reference (1550 nm C-band)

- Standard G.652/G.657 SMF: 0.180 - 0.20 dB/km
- Premium G.654 SMF: 0.154 dB/km
- Physical minimum for silica glass (Rayleigh floor): ~0.14 - 0.15 dB/km
- NANF / DNANF (volume production, 2026): 0.10 - 0.15 dB/km
- NANF / DNANF (laboratory): 0.091 dB/km lowest average loss at 1550 nm; 0.04 dB/km lowest reported minimum attenuation

In practical terms, where a signal in SMF requires an amplifier every 80 km, NANF at 0.10 dB/km would on paper support a span of roughly 160 km on the same 16 dB budget. Allowing margin for splices and for the 0.3 - 0.6 dB of SMF interfacing incurred at each amplifier site, a working figure of 120 - 140 km is realistic, reducing amplifier count by a third or more on a typical long-haul route. This translates into lower operational expenditure (OpEx) for network operators with fewer active components to power, maintain, and replace resulting in fewer active failure points along remote and submarine routes, and a meaningfully reduced carbon footprint for hyperscalers committed to Scope 2 and Scope 3 emissions reduction targets.

5. Performance Parameters

The following table provides a comparison of NANF against the two most widely deployed conventional fibre standards: G.652/657 (the global defaults) and G.654 ultra-low-loss SMF (used in premium long-haul routes).

Parameter	NANF (2026)	G.652/657	G.654	HCF advantage
Attenuation @ 1550 nm	0.091 dB/km (lab) / 0.10 - 0.15 dB/km (production)	0.18 - 0.20 dB/km	0.154 dB/km	2 - 5×
Propagation speed	99.7% of vacuum (c)	~68% of c	~68% of c	~1.5×
Latency per km	~3.34 µs/km	~4.88 µs/km	~4.88 µs/km	~30%
Chromatic dispersion	1 - 5 ps/nm/km	16 - 18 ps/nm/km	20 - 22 ps/nm/km	4 - 18×
Nonlinearity (n_2/A_eff)	~0.0001× SMF	Reference (1×)	~0.7× SMF	~10,000×
Transmission bandwidth	66 THz below 0.2 dB/km	26 THz below 0.2 dB/km	26 THz below 0.2 dB/km	~2.5×
Max power handling (nonlinear limit†)	>60 dBm (kW-class demonstrated)	~23 dBm (SBS onset; damage threshold ~36 dBm)	~23 dBm (SBS onset; damage threshold ~36 dBm)	>>100×
MFD (Mode Field Diameter)	20 - 35 µm	10.4 µm	12 - 13 µm	Larger
HCF - HCF splice loss	0.04 - 0.16 dB	<0.05 dB (SMF)	<0.05 dB	Specialist equip.
HCF - SMF interface loss	0.15 - 0.30 dB	N/A	N/A	Managed at design
Min. bend radius (install)	20 - 30 mm	7.5 mm (G.657.A2)	15 mm	More sensitive
Fibre cost (2026)	>$4,000/km	$10 - 15/km	$20 - 80/km	Premium
Standards	ITU-T draft (2027 est.)	Mature (G.652)	Mature (G.654)	Emerging

Latency: The propagation speed advantage, light at 99.7% of vacuum speed versus 68% in glass may sound abstract, but on a 100 km route it represents approximately 154 µs one-way latency saving (308 µs round-trip). On a London - Frankfurt route (~900 km of fibre on a typical low-latency path), it amounts to roughly 1.4 ms one-way (2.8 ms round-trip), worth tens of millions of pounds per year in arbitrage revenue to a large algorithmic trading firm. On AI clusters where thousands of GPU-to-GPU messages are exchanged per training step, the aggregate effect of 30% lower latency compounds significantly across the network.

Chromatic dispersion: SMF's dispersion of 16 - 18 ps/nm/km means that a pulse 1 nm wide in wavelength spreads by 18 picoseconds for every kilometre it travels. At 400G and above, this requires increasingly sophisticated digital signal processing (DSP) at the receiver to undo the spreading. NANF's low dispersion of 1 - 5 ps/nm/km enables simplified coherent receiver design, reducing DSP power consumption, and enabling higher symbol rates over the same distances.

Nonlinearity: The effective nonlinearity of NANF is approximately four orders of magnitude lower than SMF. In practical terms this means that even at very high launched powers (>30 dBm), NANF does not generate the Kerr effect nonlinearities (self-phase modulation, cross-phase modulation, four-wave mixing) that limit capacity in SMF-based DWDM systems. This has two important advantages - more channels can be packed into the transmission window without cross-channel interference, and launched power can be increased to extend unrepeatered span lengths, something not possible in SMF without generating proportionally more nonlinear noise.

6. Splicing And Connecting

Joining lengths of NANF to each other or connecting NANF to conventional single-mode fibre (SMF), is one of the most demanding field operations in modern optical network deployment. The challenge stems from a single physical fact: NANF is hollow. The thin glass walls that make it perform so well are also the walls that standard splicing equipment will instantly destroy if used without modification.

6.1 Why Standard Splicing Destroys NANF

Conventional SMF splicing works by melting both fibre ends together at well above 1,600°C. Two solid glass faces fuse into one continuous strand. This works because there is nothing inside to damage.

NANF has six hollow capillary tubes arranged around its air core, each with walls only 100 - 500 nanometres thick, roughly 200 times thinner than a human hair. When a standard arc is applied, those walls melt in milliseconds. Surface tension then pulls the molten glass outward toward the outer jacket. The hollow geometry is erased and the air core collapses into a solid plug of glass. The splice becomes a wall rather than a window and the signal cannot pass.

This failure is immediate and irreversible. There is no partial damage mode: a single incorrectly programmed arc will ruin the splice completely. This is the operational constraint that separates NANF deployment from all previous fibre installation experience.

6.2 How NANF Is Spliced Successfully

The solution is to prevent the heat from reaching the hollow core. Two approaches are used in practice, both of which keep the arc energy in the outer glass cladding and allow it to diffuse inward slowly enough to fuse the fibre surfaces without collapsing the microstructure inside.

The Ring of Fire method is the standard approach for production and field work. Specialist splicers (Fujikura FSM-100P-HCF, Sumitomo TYPE-81C-HCF) use a multi-electrode geometry that positions the fibre away from the hottest part of the plasma arc. The outer cladding is heated by indirect radiation rather than direct contact, and the heat travels inward at a controlled rate, fusing the outer surfaces before it can reach the internal structure.

Pulsed low-current discharge is a laboratory technique that uses arc currents roughly half those used for SMF (~10 mA versus ~17 mA), applied in short bursts. Each burst adds a small increment of heat. The cumulative energy softens the outer glass just enough to bond the fibres without triggering thermal runaway inward. This approach requires more time per splice and is not suited to field deployment at scale but is valuable for prototype assembly and research environments.

Both methods are NANF-specific. They cannot be approximated on a standard SMF splicer by adjusting current alone. The arc programme geometry is fundamentally different, and a splicer without a certified HCF profile cannot be used with NANF.

6.3 Connecting NANF to Standard SMF

A NANF network currently needs at least two junctions where it connects to conventional SMF, at the transmitter and at the receiver. These junctions require careful engineering because the two fibre types guide light in fundamentally different ways.

The core problem is mode field size mismatch

Standard SMF carries its signal in a mode-field beam roughly 10 micrometres in diameter. NANF, which needs a larger air core to achieve its ultra-low loss, carries its signal in a beam 20 - 35 micrometres across. Joining them directly is like connecting a garden hose to a drainpipe, most of the light misses the target entirely, resulting in 2 - 3 dB of loss, immediately cancelling the attenuation advantage that makes NANF attractive in the first place.

Three established methods bridge this gap, each with different trade-offs:

- **GRIN lens:** A short segment of graded-index fibre spliced to the SMF end acts as a miniature lens, expanding and reshaping the SMF beam to match the NANF mode before they meet. This is the most widely deployed commercial approach.

- **Thermally Expanded Core (TEC):** The SMF tip is heated slowly over an extended period, causing its germanium dopant to diffuse outward and widen the guided beam. The expanded SMF beam then matches the NANF profile more closely before the final fusion step.

- **Solid end-cap:** The final 50 - 100 µm of the NANF tip is deliberately melted into a solid silica block before the junction is made. The SMF is then fused to this solid stub in a standard splice. This creates a sealed, permanently hermetic interface that isolates the hollow core from moisture and contamination, making it the preferred approach for subsea cables and any harsh outdoor environment.

The fixed 0.3 - 0.6 dB overhead from the two SMF interfaces per HCF segment is worth understanding strategically: it favours long, continuous NANF runs where the overhead is a small fraction of the total link budget. On very short routes with many patching points, the accumulated interface penalties progressively erode the attenuation advantage.

6.4 Testing and Verification in the Field

Two testing assumptions that are standard practice on SMF do not transfer to NANF:

- **Splicer loss estimates are not reliable.** Standard field splicers estimate splice quality by imaging the core alignment optically. Because the hollow NANF core and its sub-micron capillary structure cannot be resolved through the outer cladding, the splicer's on-screen loss reading is not meaningful for NANF joints. Every NANF splice must be verified by bidirectional OTDR measurement.

- **Standard OTDR distance readings will be wrong on hybrid links.** Light travels through NANF at 99.7% of vacuum speed, compared with ~68% in SMF. An OTDR configured for SMF will calculate distances based on the slower propagation speed and will therefore mislocate events on the NANF portion of a hybrid link by a consistent offset. Field testers must use equipment with a dedicated HCF refractive index profile for any measurement on a hybrid route.

7. Commercial Opportunities

NANF is the next-generation default for every optical link where latency, capacity, or power handling matter more than initial fibre cost. As manufacturing costs fall from thousands of dollars per kilometre today toward parity with standard SMF by 2035, NANF will capture an expanding share of the world's network infrastructure. The commercial question is which markets will adopt it first, which use cases will generate the highest returns, and where manufacturers are positioned to capture that demand.

7.1 Where the Commercial Case Is Already Proven

Two use cases have moved beyond pilot and are generating revenue today. Both convert the latency advantage of NANF into economic value.

High-frequency trading routes represent the clearest ROI case in optical networking. Every 100 km of HCF saves 154 µs of one-way latency (308 µs round-trip) compared with SMF.

AI and hyperscale data centre interconnects represent the largest volume opportunity. GPU clusters for large language model training exchange vast numbers of messages per training step, and the synchronisation latency between GPUs directly constrains training throughput.

7.2 Sectors at the Threshold of Adoption

Several further sectors are approaching commercial deployment towards 2030, as manufacturing costs fall and the supporting ecosystem matures.

Long-haul backbone and subsea cables benefit from two compounding advantages. NANF's lower attenuation extends unamplified terrestrial spans from 80 km to 120 km under the same loss budget, reducing active amplifier count by approximately 33% on a typical route. Fewer repeaters means lower capital cost, lower operational expenditure, fewer failure points on remote routes, and a meaningfully smaller carbon footprint per terabit transmitted.

5G and 6G fronthaul is a volume opportunity with a specific technical driver. Time-division duplex 5G imposes strict latency budgets on the connection between baseband processing units and remote radio heads. Conventional SMF limits this separation to roughly 20 km. HCF's ~30% latency advantage extends the viable separation to approximately 29 km (20 km × 1.46), allowing a single baseband unit to serve a wider geographic area. In dense urban markets where siting baseband equipment is expensive and constrained, this converts to a direct infrastructure cost reduction.

Quantum communications use the suppression of spontaneous Raman scattering from co-propagating classical channels into the quantum band. NANF could become the preferred physical medium for quantum key distribution, with BT having demonstrated a 6 km NANF link at 1.1 Mbps secret-key rate, with the real HCF benefit being crosstalk suppression allowing quantum and classical channels to share a fibre. As national quantum network programmes accelerate across the UK, EU, Japan, and India, NANF could be the infrastructure of choice for the physical layer.

7.3 Geographic Markets: Where to Deploy and When

NANF deployment is not evenly distributed. The markets where it is taking hold first share a common set of characteristics: dense concentrations of financial or AI infrastructure, established operator appetite for premium performance, and procurement organisations capable of evaluating a total-cost-of-ownership case rather than a per-metre price comparison.

North America and Western Europe

These are the most developed markets. Microsoft's Azure deployment and the euNetworks financial trading routes have established proof points at production scale. Western Europe also offers a growing market in quantum network infrastructure through funded national programmes in the UK, Germany, France, and the Netherlands.

India

Three converging forces are creating demand simultaneously: the AI data centre build-out across Mumbai, Delhi-NCR, Hyderabad, and Chennai; the emergence of GIFT City as an internationally competitive financial hub seeking to attract high-frequency trading infrastructure; and the 5G rollout across dense metro corridors where fronthaul latency budgets are already being pushed.

Middle East and GCC

The Gulf Cooperation Council states offer a structurally attractive market for a different reason: greenfield advantage. With large-scale infrastructure programmes underway, NANF can be designed in from the outset rather than retrofitted into existing SMF networks. Sovereign AI investment mandates under Vision 2030 and the UAE AI Strategy 2031 are driving hyperscale data centre demand at a pace where performance differentiation justifies cost premium.

Southeast Asia

Singapore's density of hyperscale campuses, combined with the concentration of subsea cable landing stations creates a compelling two-part opportunity. The terrestrial segments connecting landing stations to Singapore's major data centre campuses are short, high-value routes well-suited to NANF. The regional subsea cables from Singapore to Indonesia and Malaysia, along with the emerging India - ASEAN routes, are longer, and the long unamplified spans translate directly into lower capital cost and reduced long-term maintenance exposure on remote submarine routes.

8. Conclusion

Hollow core fibre has crossed the line from laboratory research to commercial infrastructure. That transition, which took conventional SMF decades has compressed into a few years for NANF. Driven by the convergence of record-breaking loss figures, AI & hyperscaler deployments at scale, and a rapidly maturing supply and testing ecosystem. The question for operators and manufacturers is no longer whether NANF will achieve broad deployment, but how to position for the market as it develops.

NANF guides light at 99.7% of vacuum speed through an air core, eliminating Rayleigh scattering and achieving attenuations that silica-core fibre cannot reach by any amount of material or process refinement. Attenuation of 0.04 dB/km points toward a new performance level, and modelling suggests that 0.01 dB/km or better may eventually be reachable. Surface scattering, the dominant loss mechanism in current NANF is an engineering problem being addressed steadily through tighter drawing control, in-situ fabrication monitoring, and improved preform geometry.

The commercial case is proven in the use cases that matter most. High-frequency trading routes generate premium recurring revenue with sub-one-year payback. Hyperscale AI interconnects are already deploying at scale: Microsoft's 1,200+ km of live Azure DNANF confirms that laboratory performance can be manufactured and maintained in production networks. Long-haul backbone and subsea routes follow the economic principle of fewer amplifiers, which translates into lower capital cost, lower operational expenditure, and a smaller carbon footprint per terabit transmitted.

The remaining barriers are engineering and ecosystem challenges. Amplification for long-haul NANF is maturing through hybrid HCF/SCF EDFA architectures. Splicing has improved from a specialist laboratory operation to a repeatable field technique. Testing and certification tools are reaching the field. Standards are in draft. Cost is on a steep learning curve toward parity with premium SMF within a decade.

NANF is not a technology that will gradually supplement conventional fibre over an extended transition period. NANF will become the standard on routes where latency, capacity, or power handling justify the cost, a category growing rapidly as AI infrastructure scales and financial networks intensify. The manufacturers and operators who establish positions now will define the reference architecture for the next generation of global optical infrastructure.

References

[1] F. Poletti, "Nested antiresonant nodeless hollow core fiber," Optics Express.

[2] M. Zeisberger and M. A. Schmidt, "Analytic model for the complex effective index of the leaky modes of tube-type anti-resonant hollow core fibers," Scientific Reports.

[3] E. Numkam Fokoua, S. Abokhamis Mousavi, G. T. Jasion, D. J. Richardson and F. Poletti, "Loss in hollow-core optical fibers: mechanisms, scaling rules, and limits," Advances in Optics and Photonics.

[4] W. Belardi and J. C. Knight, "Hollow antiresonant fibers with reduced attenuation," Optics Letters.

[5] B. Debord et al., "Ultralow transmission loss in inhibited-coupling guiding hollow fibers," Optica.

[6] A. Taranta et al., "Exceptional polarization purity in antiresonant hollow-core optical fibres," Nature Photonics.

[7] B. J. Mangan et al., "Low loss (1.7 dB/km) hollow core photonic bandgap fiber," OFC Conference 2004.

[8] F. Poletti, M. N. Petrovich and D. J. Richardson, "Hollow-core photonic bandgap fibers: technology and applications," Nanophotonics.

[9] T. D. Bradley et al., "Record low-loss 1.3 dB/km data transmitting antiresonant hollow core fibre," ECOC 2018.

[10] T. D. Bradley et al., "Antiresonant hollow core fibre with 0.65 dB/km attenuation across the C and L telecommunication bands," ECOC 2019.

[11] G. T. Jasion et al., "Hollow core NANF with 0.28 dB/km attenuation in the C and L bands," OFC 2020.

[12] H. Sakr et al., "Hollow core NANFs with five nested tubes and record low loss at 850, 1060, 1300 and 1625 nm," OFC 2021.

[13] G. T. Jasion et al., "0.174 dB/km hollow core double nested antiresonant nodeless fiber (DNANF)," OFC 2022.

[14] Y. Chen et al., "Hollow core DNANF optical fiber with <0.11 dB/km loss," OFC 2024.

[15] M. N. Petrovich, E. Numkam Fokoua, Y. Chen et al., "Broadband optical fibre with an attenuation lower than 0.1 decibel per kilometre," Nature Photonics.

[16] G. Amouzad Mahdiraji et al., "First dual-band hybrid window antiresonant HCF with 0.13 dB/km loss at 1 μm and 0.11 dB/km at 1.55 μm," OFC 2026.

[17] R. Slavík, G. Marra, E. Numkam Fokoua, N. Baddela, N. V. Wheeler, M. Petrovich, F. Poletti and D. J. Richardson, "Ultralow thermal sensitivity of phase and propagation delay in hollow core optical fibres," Scientific Reports.

[18] ITU-T Study Group 15, Executive Summary, October 2025.

[19] "Hollow-core fiber in direct-detection optical networks: technology readiness, deployment drivers, and adoption outlook," arXiv 2026.

[20] BT, "BT conducts world's first trial of quantum-secure communications over hollow core fibre cable," press release, 10 September 2021.

[22] YOFC, "YOFC unveils game-changing hollow-core fibre advances at OFC 2025," April 2025.

[23] YOFC, "YOFC unveils HollowBand hollow-core fibre at MWC Barcelona 2026," March 2026.

[24] YOFC, "OFC 2026: Luo Jie 'Hollow core fibres: from laboratory innovations to commercial applications'," March 2026.

HFCL

About HFCL

HFCL is a leading technology company specialising in creating digital networks for telcos, enterprises, and governments. Over the years, HFCL has emerged as a trusted partner offering sustainable high-tech solutions with a commitment to providing the latest technology products to its customers. Our strong R&D expertise, coupled with our global system integration services and decades of experience in fibre optics, enable us to deliver innovative digital network solutions required for the most advanced networks.

The Company's in-house R&D Centres located at Gurgaon & Bengaluru, along with invested R&D houses and other R&D collaborators at different locations in India and abroad, innovate a futuristic range of technology products and solutions. HFCL has developed capabilities to provide premium quality Optical Fibre and Optical Fibre Cables, state-of-the-art telecom products including 5G Radio Access Network (RAN) products, 5G Transport Products, Wi-Fi Systems (Wi-Fi 6, Wi-Fi 7), Unlicensed Band Radios, Switches, Routers, and Software Defined Radios.

The Company has state-of-the-art Optical Fibre and Optical Fibre Cable manufacturing plants at Hyderabad, an Optical Fibre Cable manufacturing plant in Goa, and its subsidiary, HTL Limited, at Chennai.