Forwards or Backwards

https://forwardorbackwards.com



The Undersea Cables: The Web's 99% Vulnerability

Watch on YouTube: https://www.youtube.com/watch?v=mN4nfe6-4b0

Chapters

1. A thread on the seabed
2. Why it is not the sky
3. What is actually inside
4. The repeaters and the power
5. Choosing a route
6. The shape of the map
7. Who owns them now
8. Landing stations

9. Two hundred faults a year
10. The ships that mend them
11. February 2024: the Red Sea
12. What the Red Sea actually showed
13. October 2023: the Balticconnector
14. November 2024: the Yi Peng 3
15. December 2024: the Eagle S
16. The difficulty of proving intent
17. Baltic Sentry
18. Redundancy is the real protection
19. Rules, and who makes them
20. The thread, and what rests on it

1. A thread on the seabed



Somewhere off the coast of Cornwall, under a few metres of silt, lies a cable about as thick as a garden hose. Inside it sit a handful of glass fibres, each roughly the width of a human hair. Through those fibres passes a substantial share of everything Britain says to the rest of the world.

Not a metaphor for the internet. The internet. Your bank transfer, your video call, the instruction that settles a trade in New York, the message that reaches a relative in Lagos. All of it, as light, inside glass, on the floor of the sea.

There is a comfortable assumption that the modern world runs on something abstract. The cloud. The network. The air. It does not. It runs on a few hundred physical objects lying in the dark, most of them unguarded, most of them unmarked, and almost all of them owned by private companies.

As of early 2026, TeleGeography tracks more than six hundred active and planned submarine cables, and believes there are over one and a half million kilometres of submarine cable in service worldwide. That is enough to circle the equator nearly forty times.

This film is about what those cables actually are, how they break, who mends them, and what happens when several break at once in the same place. It is also about a question that has moved, in the last three years, from a technical curiosity to a matter of national defence.

That question is simple. If the thing everything depends on is a thread on the seabed, how much protection does a thread have?

The answer turns out to be less than you would expect, and more than you might fear. Both halves of that sentence matter.

2. Why it is not the sky



Most people, asked where international internet traffic goes, would guess satellites. It is an understandable guess. Satellites are visible, they are discussed, and they carry the association of modernity.

The numbers do not support it. According to TeleGeography, satellites account for zero point three seven per cent of all United States international capacity. Not thirty-seven per cent. Not three point seven. Zero point three seven.

There are two reasons, and neither is about ambition.

The first is physics. A geostationary satellite sits roughly thirty-six thousand kilometres above the equator. A signal must travel up and back down, and light, fast as it is, takes time. That round trip adds something in the region of half a second to a conversation. For a phone call it is awkward. For a financial trade, a video call, or a remote surgical instrument, it is unusable.

Low earth orbit constellations reduce that delay considerably, because they fly much closer. They have genuinely changed what is possible in places a cable will never reach. But they do not change the second reason.

The second reason is capacity. A single modern submarine cable can carry hundreds of terabits per second. An entire satellite constellation, shared across every user beneath it, does not approach what one cable delivers into one landing station. The economics are not close either. A cable's capacity is

dedicated, and its cost per bit falls as the equipment on shore is upgraded without anyone touching the cable itself.

So satellites complement the cables. They do not replace them, and on current physics they will not. When someone says the internet is in the cloud, the honest correction is that the cloud is mostly underwater.

3. What is actually inside



Hold a section of deep-water submarine cable and the first surprise is how thin it is. On the abyssal plain, where nothing is likely to touch it, a cable can be around seventeen to twenty millimetres across. Comparable to a garden hose.

Work inward from the outside. There is a polyethylene sheath. Beneath it, a copper or aluminium conductor that carries electrical power, not data. Beneath that, layers of polycarbonate, an aluminium water barrier, petroleum jelly, and a steel tube. Inside the steel tube, suspended in more jelly, sit the optical fibres. Often no more than four to sixteen of them.

Everything outside those fibres exists to protect them and to feed them power. The data itself occupies almost none of the cross-section.

Closer to shore the cable changes character entirely. In shallow water, where anchors drop and trawls drag, the same cable is wrapped in one or two layers of galvanised steel armour wire and can reach the thickness of a human arm. It is frequently buried by a sea plough that cuts a furrow, lays the cable in, and lets the seabed close over it.

This is a deliberate engineering judgement about risk. Deep water is statistically safe, so the cable is light and cheap. Shallow water is dangerous, so the cable is heavy and expensive. The great majority of faults happen in water shallower than two hundred metres, which is precisely where human activity is.

The vulnerability, in other words, is not evenly spread. It is concentrated exactly where the coastline is.

4. The repeaters and the power



Light dims as it travels. Over thousands of kilometres of glass, a signal launched from Cornwall would arrive in New York as nothing at all. So every sixty to eighty kilometres, a transatlantic cable carries a repeater: a pressure-resistant housing, usually a metre or so long, containing an optical amplifier.

These are not repeaters in the sense of receiving and retransmitting. They are erbium-doped fibre amplifiers. A short length of fibre is doped with the rare-earth element erbium, and a pump laser excites it. When the weakening signal passes through the excited section it stimulates emission at the same wavelength, and the light emerges stronger than it went in. The signal is never converted back into electricity. It is amplified while remaining light.

That amplification needs power, and here is the part that surprises people. The power comes from the shore. A power feed equipment room at each landing station puts a high direct-current voltage — thousands of volts — onto the copper conductor, and every repeater along the cable taps a little of it in series. A single cable can be energised from both ends at once, with the mid-point floating near zero.

It means a transatlantic cable is a live electrical circuit thousands of kilometres long, and an object nobody can reach without a ship.

It also means a repeater is a component that must work, untouched, in eight thousand metres of water, at pressures around eight hundred atmospheres, for a design life of twenty-five years. Not maintained. Not inspected. Simply working, in the dark, for a quarter of a century.

5. Choosing a route



Before a cable is laid, it is surveyed. A survey vessel spends weeks or months running the proposed corridor, mapping the seabed with multibeam sonar and sub-bottom profilers, sampling sediment, recording every obstacle.

The route that results is never the shortest line. It is a negotiated path that avoids submarine canyons, seamounts, known landslide zones, active fishing grounds, anchorages, dumped munitions, other cables, pipelines and marine protected areas. It threads through the territorial waters and exclusive economic zones of every country it passes, each of which may require its own permit.

The engineering aim is simple to state and hard to achieve. Put the cable where nothing will disturb it, and where, if something does, a ship can reach it.

Laying follows the survey. A cable ship pays out the cable over a stern sheave at a few knots, matching payout speed to ship speed and water depth so the cable reaches the seabed with the right slack. Too little slack and it hangs suspended across a hollow, vibrating in the current until it fatigues. Too much and it loops, snags and costs money.

In shallow water a plough is towed behind, burying the cable as it goes, typically one to three metres down.

The whole operation is closer to agriculture than to telecoms. It is slow, weather-dependent and governed by the sea. A transatlantic system takes weeks to lay, and often years to permit.

6. The shape of the map



Look at a submarine cable map and the first thing you notice is that it is not a web. It is a set of bundles.

Cables do not spread evenly across the ocean. They cluster into corridors, because the physics of routing, the economics of landing stations and the geography of continents all push in the same direction. The North Atlantic carries a dense bundle between the eastern United States and western Europe. Another runs from Europe through the Mediterranean, down the Red Sea, and on to India and South East Asia. A third crosses the Pacific.

Within those corridors, the cables converge further at specific points. The Strait of Malacca. The Suez and Red Sea passage. The Luzon Strait north of the Philippines. The approaches to Marseille, to Alexandria, to Fujairah, to Singapore. In these places, cables that are thousands of kilometres apart in mid-ocean run within sight of one another.

That convergence is not a mistake. It is the direct consequence of geography: there is only one Red Sea, and everything from Europe to Asia either goes through it or goes the long way round Africa, adding thousands of kilometres and milliseconds that customers will not pay for.

The result is a network that is genuinely resilient in the open ocean and genuinely concentrated at the pinch points. A single break in the middle of the Atlantic is an inconvenience absorbed by rerouting. Several breaks in the same narrow channel are something else.

This is the single most important fact about the physical internet, and it is the one least reflected in how people imagine it. The map looks like a mesh. The chokepoints behave like a funnel.

7. Who owns them now



For most of the history of submarine cables, they were built by consortia of telecommunications carriers. A dozen national operators would each take a share, each get capacity in proportion, and each carry everyone else's traffic under reciprocal agreements. The cable was shared infrastructure, built by the companies whose business was carrying other people's messages.

That has changed, and quickly. Over the past decade the largest content companies — the businesses that run search, video, social networks and cloud platforms — have become the dominant investors in new submarine capacity. They now own or hold major stakes in a substantial share of newly built transoceanic systems, and in several cases build cables alone, for their own traffic.

The logic is straightforward. If your business depends on moving enormous volumes of data between your own data centres on different continents, buying capacity from a carrier is a cost you do not control. Building the cable turns an operating expense into an asset, and gives you certainty about capacity and routing.

The consequences are less straightforward. Infrastructure that was once collectively owned by regulated national carriers is now substantially owned by a small number of private firms answerable to shareholders in a handful of jurisdictions. Those firms have been, by most accounts, good stewards — they invest heavily, they build redundancy, and their commercial interest in uptime is absolute.

But the question of who decides where the world's connectivity goes, and which regions get a new cable and which wait, has moved. It now sits largely with companies, not governments. That is a

change in the character of critical infrastructure, and it happened without much public discussion.

8. Landing stations



A cable has to come ashore, and where it does is one of the least glamorous and most sensitive parts of the whole system.

A cable landing station is usually a plain, low building set back from a beach. It has no signage worth the name. From the outside it could be a pumping station or a small telephone exchange. Inside are the power feed equipment that energises the cable, the line terminal equipment that turns light into routable traffic, batteries, generators, and the backhaul that carries everything inland to the nearest major network node.

Each station typically serves several cables. Each cable typically lands at only two or three places along its length. The set of landing points for an entire country can often be counted on two hands, and in smaller nations on one.

This is the concentration problem in its sharpest form. The cables spread out across an ocean, but they come together at the shore, and then they come together again in the backhaul. In some countries, several cables that took entirely different routes across the sea terminate within a few kilometres of each other and then share the same terrestrial fibre inland.

A network can be perfectly redundant at sea and still have a single point of failure on land. Regulators have begun to notice this. Operators have known it for years. Fixing it is expensive, unglamorous, and produces no new capacity — which is precisely why it is so often deferred.

9. Two hundred faults a year



Cables break constantly. This is the fact that most surprises people, and it is the fact that should reassure them most.

On average there are roughly two hundred faults each year, according to data from the International Cable Protection Committee. That is around four every week, somewhere in the world, every week of the year.

Almost none of them are sinister. Accidents involving fishing vessels and ships dragging anchors account for two-thirds of all cable faults. The rest are largely natural: submarine landslides, turbidity currents, earthquakes, abrasion where a cable rests on rock, and simple component failure at the end of a long design life.

The pattern follows human activity precisely. Faults cluster in shallow water, near ports, on fishing grounds, and in the approaches to busy straits. A cable in four thousand metres of water, a hundred miles from anywhere, is very rarely troubled.

Two hundred faults a year, and almost nobody notices. That is not luck. It is the direct result of building a network with enough spare capacity and enough alternative paths that a single break can be routed around automatically, often within seconds, while a repair ship is still being tasked.

The resilience is real. It is engineered, it is paid for, and it works. The question that the last three years have raised is not whether the system copes with random accidents — it plainly does — but whether it copes with several deliberate ones, in the same place, at the same time.

10. The ships that mend them



Repair is slow, physical, and carried out by a fleet most people have never heard of.

When a fault occurs, the operator locates it electrically from shore, often to within a few hundred metres, by measuring how a test pulse reflects back. A cable ship is then tasked, loads spare cable and a jointing team, and sails.

At the site the ship drags a grapnel along the seabed until it catches the cable, cuts it, and recovers one end to the surface. That end is buoyed off. The ship then recovers the other end, splices in a fresh section of cable long enough to bridge the gap plus the slack needed to reach the surface twice, tests the joint, and lowers the repaired bight back to the bottom.

In good weather, in accessible water, this takes days. In deep water, in poor weather, or where permits are required from a coastal state before work can begin, it takes weeks.

And the fleet that does it is thin. A 2025 assessment by TeleGeography and Infra-Analytics put the investment needed at roughly three billion dollars to sustain current service levels and avoid repair delays, entailing the acquisition of fifteen replacement ships and five additional ones — twenty vessels in total.

That figure is worth sitting with. The entire physical internet depends on a global fleet of specialist ships small enough that adding twenty would transform it, and old enough that fifteen need replacing. There is no surge capacity. If several cables break at once in different oceans, they are repaired in sequence, because there is no one else to send.

11. February 2024: the Red Sea



On the twenty-fourth of February 2024, three submarine cables were cut in the Red Sea: Seacom, which also carries TGN-Eurasia traffic, the Europe India Gateway, and Asia-Africa-Europe-1.

The leading explanation involves a ship called the Rubymar. Days earlier it had been struck by missiles fired from Yemen. According to analysis by the network intelligence firm Kentik, the crew dropped anchor and abandoned the crippled vessel, and the Rubymar then began to drift, dragging its anchor across the seabed. The ship sank on the second of March.

An anchor dragged across the floor of the Red Sea is, as Kentik put it, one of the top causes of submarine cable cuts. It is worth being precise about the sequence, because it matters: the cables were almost certainly not the target. They were collateral damage from a vessel that had been attacked for other reasons and then left to drift.

The effect was measurable across two continents. Kentik's traffic analysis showed disruption in East Africa — Tanzania, Kenya, Uganda, Mozambique — and as far away as South East Asia, with observable shifts in transit paths for Vietnamese and Emirati networks. The Seacom disruption registered at nine forty-six UTC, Asia-Africa-Europe-1 five minutes later.

Nobody lost the internet. Traffic rerouted, as it is designed to. But it rerouted onto longer paths, at higher cost, with worse latency, and with less headroom for the next fault.

And the repair was not quick. Working in the Red Sea in 2024 meant working in a war zone, requiring permits and security assurances before a cable ship could sit stationary over a fault for days at a time.

12. What the Red Sea actually showed



The Red Sea cuts were revealing not because of the damage but because of what the damage exposed.

Three cables in one narrow channel went down in a single incident. They were operated by different companies, served different customers, and had been laid at different times. On a map they are distinct systems offering genuine diversity. On the seabed they were close enough that one dragging anchor could take all three.

This is the gap between logical redundancy and physical redundancy, and it is the central problem in cable resilience. A customer buying capacity on two different cables is entitled to assume that a fault on one will not affect the other. That assumption holds in the open ocean. It does not necessarily hold in a strait that everything must pass through.

The industry term is shared risk link group — the set of circuits that would fail together because they share a physical hazard. Identifying those groups requires knowing exactly where every cable lies, which is commercially sensitive information that operators do not freely share with each other or with customers.

So the buyer of diverse capacity often cannot verify the diversity they are paying for, and the operator selling it may not have full visibility either.

The honest summary is this. The system is far more robust than the alarming headlines suggest, and far less diverse at the chokepoints than the reassuring maps imply. Both things are true at the same

time, and which one matters depends entirely on where the damage happens.

13. October 2023: the Balticconnector



On the eighth of October 2023, the Balticconnector gas pipeline between Finland and Estonia was damaged, along with a telecommunications cable in the same area. Finnish investigators subsequently linked the damage to a vessel called the Newnew Polar Bear, which had been in the area, and to an anchor recovered from the seabed near the damage site.

At the time, the incident was treated as unusual. In hindsight it reads as the beginning of a pattern.

What made it different from the Red Sea was not the mechanism – it was again an anchor dragged along the bottom – but the context. This was not a ship that had been attacked and left to drift. It was a vessel under way, in peacetime, in a shallow sea surrounded by NATO members, that appeared to have dragged its anchor for a considerable distance across infrastructure belonging to several countries.

The Baltic is a particular case. It is shallow, it is busy, and it is enclosed. Its average depth is a little over fifty metres – well inside the band where cables are most at risk – and it is crossed by a dense network of power interconnectors, gas pipelines and data cables linking countries with strained relations.

It is, in short, the place where an anchor does the most damage for the least effort, and where the question of intent is hardest to separate from the question of incompetence.

That question would be asked repeatedly over the following fifteen months.

14. November 2024: the Yi Peng 3



On the seventeenth of November 2024, at around ten in the morning local time, the BCS East-West Interlink was severed. It is a two hundred and eighteen kilometre cable running between Šventoji in Lithuania and Katthammarsvik on the Swedish island of Gotland, and it carried approximately one third of Lithuania's internet capacity.

Some sixteen hours later, at two in the morning UTC on the eighteenth, the C-Lion1 cable between Finland and Germany failed off the Swedish island of Öland. Its operator, Cinia, stated that the cable had been severed by an outside force. The second fault lay roughly a hundred kilometres from the first.

Suspicion fell on a Chinese-flagged bulk carrier, the Yi Peng 3, which had passed over both locations in the intervening hours. Reporting by the Wall Street Journal alleged that the ship dropped one of its anchors while its engines continued to drive the vessel forward, ploughing through the seabed over a distance of more than a hundred miles.

Not everyone accepted that account. Writing in the Daily Telegraph, the naval commentator Tom Sharpe argued that a normal anchor would be too heavy to drag at the vessel's reported seven knots, and raised the possibility of falsified AIS position data or some other explanation.

Both cables were restored on the twenty-eighth of November 2024. The ship was shadowed by European navies, and eventually sailed. No prosecution followed.

15. December 2024: the Eagle S



Seven weeks later, on Christmas Day 2024, it happened again and on a larger scale.

A tanker named the Eagle S, sailing in the Gulf of Finland, was linked to damage affecting the Estlink 2 power interconnector between Finland and Estonia and four telecommunications cables in the same waters. One incident, five pieces of critical infrastructure, on the same day.

Estlink 2 matters beyond data. It is an electricity interconnector, and its loss removed a substantial share of the transfer capacity between the Finnish and Estonian power systems at the start of winter.

This time the response was different. Finnish authorities boarded the vessel, escorted it into Finnish waters, and detained it. Finnish investigators said publicly that the damage appeared to have been intentional. It was the first time in the sequence that a state had physically taken control of a suspect ship and pursued a criminal case.

The Eagle S belonged to what is commonly described as a shadow fleet: ageing tankers with opaque ownership, unclear insurance and flags of convenience, operating largely outside the reach of Western regulation. Whether their crews are agents of anything, or simply badly run ships making bad decisions, is exactly the ambiguity that makes the pattern so difficult to address.

In early 2025 the sequence continued. On the third of February 2025 Swedish authorities determined that a cable break involving the vessel Vezhen, on a link between Sweden and Latvia, was accidental rather than deliberate.

16. The difficulty of proving intent



Here is the uncomfortable part, and the reason this subject resists confident conclusions.

For the November 2024 Baltic cuts, United States officials assessed that the cables had not been cut deliberately. For the Christmas Day incident, Finnish investigators said the damage appeared intentional. For the February 2025 Vezhen break, Swedish authorities concluded it was an accident.

Same sea. Same mechanism. Three different official conclusions, reached by three capable states with access to intelligence the public does not have.

That divergence is not incompetence. It reflects a genuine evidentiary problem. A dragged anchor leaves an identical scar on the seabed whether the hand on the windlass was careless or deliberate. Proving intent requires evidence about what happened on the bridge, and that evidence is aboard a foreign-flagged ship with a foreign crew that has usually left the area by the time anyone has confirmed the damage.

The legal position compounds it. Outside territorial waters, a coastal state's authority over a passing vessel is limited. The 1884 Convention for the Protection of Submarine Telegraph Cables and the United Nations Convention on the Law of the Sea both require states to make cable damage a punishable offence — but the state obliged to prosecute is generally the flag state, which may have little interest in doing so.

So the honest position is this. There is a pattern of incidents concentrated in one strategically tense sea over a short period, which is itself a fact worth taking seriously. But for any individual incident, the

public evidence rarely supports a confident claim of sabotage, and several have been formally assessed as accidents.

Uncertainty is not the same as innocence. It is also not the same as proof.

17. Baltic Sentry



In January 2025, NATO responded. The alliance's Baltic littoral states announced an activity called Baltic Sentry, adding vessels, patrol aircraft and staff to the existing Commander Task Force Baltic, headquartered at Rostock.

The concept is straightforward and old-fashioned: presence. Put ships and aircraft where the infrastructure is, make it obvious that vessels behaving strangely will be watched, identified and approached, and change the calculation for anyone considering dragging an anchor across a pipeline.

According to an Atlantic Council assessment, no suspicious undersea incidents occurred in the Baltic Sea after January 2025. One official quoted in that assessment described it as a sign that the deterrence being provided was working.

That claim deserves care. An absence of incidents after an intervention is consistent with the intervention working. It is also consistent with the incidents having been accidents that simply stopped happening, or with the activity moving elsewhere, or with a change in intent for reasons unconnected to NATO patrols. You cannot prove a deterrent worked by pointing at events that did not occur.

What can be said with more confidence is narrower and still meaningful. The response demonstrated that a group of states could co-ordinate surveillance of seabed infrastructure at scale and at short notice, which they had not previously done. Whatever it deterred, it built a capability that did not exist before.

And it established a precedent: that damage to submarine infrastructure is now treated as a defence matter, not solely a commercial one.

18. Redundancy is the real protection



If you wanted to make the physical internet safer, patrolling it would not be the first thing you did. Nor would armouring the cables.

The protection that actually works is redundancy, and it works because it makes any single cable unimportant.

A network with genuine path diversity absorbs a break without the user noticing. Traffic reroutes automatically in seconds. The engineering that makes this possible is not exotic: it is spare capacity on alternative routes, held in reserve and paid for, plus routing systems designed to fail over without human intervention.

The cost is real. Spare capacity is capacity you have bought and are not selling. Every pound spent on a path you hope never to use is a pound not spent on growth. The commercial pressure is constantly towards running the network closer to full, on fewer paths, at higher utilisation.

The places that suffer most from cable faults are, almost without exception, the places with the fewest cables. An island nation served by one cable has no redundancy to fall back on and nothing to reroute to. When it breaks, the country goes dark and stays dark until a ship arrives, which in some regions means weeks.

This is why resilience is fundamentally a distribution problem rather than a security problem. The countries with twelve cables and four landing stations are effectively immune to single faults. The

countries with one cable are not. Adding a second cable to a nation that has one does more for global resilience than armouring every cable that already exists.

19. Rules, and who makes them



The governance of submarine cables is, to put it plainly, thin.

The foundational treaty dates from 1884. The Law of the Sea Convention adds obligations to protect cables and to punish damage, but leaves enforcement largely to flag states. There is no international body with authority to inspect a suspect vessel, to compel route disclosure, or to require a coastal state to grant a repair permit quickly.

That last point is not trivial. A cable ship cannot begin work in another country's waters without permission, and permits can take weeks. A fault that takes three days to repair can take a month to be allowed to repair.

There has been movement. In late 2024 the International Telecommunication Union and the International Cable Protection Committee established an International Advisory Body for Submarine Cable Resilience, bringing together operators, regulators and governments. In February 2025 a summit in Abuja produced a declaration on cable resilience, with commitments on faster permitting, route diversity and information sharing.

These are the right instruments. They are also voluntary, advisory, and slow.

The gap between the pace of that process and the pace of events in the Baltic is the honest measure of where the world stands. Between October 2023 and February 2025 there were repeated incidents affecting power interconnectors and data cables across a sea surrounded by wealthy, well-governed, technically capable states. The institutional response has been a declaration and an advisory body.

That is not nothing. It is not proportionate either.

20. The thread, and what rests on it



Return to Cornwall, and the cable under the silt.

It is still there. It is still about as thick as a garden hose. It is still carrying a substantial share of everything Britain says to the rest of the world, and it is still unguarded, unmarked, and owned by a company rather than a country.

Nothing in the last three years has changed the fundamental architecture. More than six hundred cables. Over one and a half million kilometres. Roughly two hundred faults a year, two-thirds of them from fishing gear and anchors. A repair fleet that needs twenty more ships and three billion dollars to hold its current service levels. A body of international law written before the telephone was widespread.

What has changed is the assumption underneath it all. The system was designed against accidents — against a trawler on a bad day, an earthquake, a slow decay of components. It was designed on the premise that nobody would want to break it on purpose, because everybody needs it.

That premise is now being tested, and the tests are ambiguous by design. An anchor is not a weapon. A drifting ship is not an attack. A pattern is not a proof.

The genuinely reassuring fact is that the network absorbed all of it. Through the Red Sea, through the Baltic, through every one of those two hundred annual faults, the internet kept working for almost everyone, almost all the time, because engineers spent decades building in more paths than anyone needed.

The genuinely uncomfortable fact is that this resilience is thinnest exactly where the pressure is greatest: at the chokepoints, at the landing stations, and in the countries with only one cable.

The thread holds. It is worth knowing how thin it is.

Sources and credits

Primary sources

- TeleGeography, *Submarine Cable Frequently Asked Questions* and the Submarine Cable Map (accessed early 2026) – more than 600 active and planned cables, over 1.5 million km in service, satellites at 0.37% of United States international capacity.
- TeleGeography and Infra-Analytics, *submarine cable repair fleet assessment* (2025) – about $3bn of investment needed to sustain current service levels, 15 replacement vessels and 5 additional ones, 20 in total.
- Cinia (operator statement, 18 November 2024) – the C-Lion1 Finland-Germany cable severed by outside force off Öland, roughly 100 km from the BCS East-West Interlink fault 16 hours earlier.
- NATO – *Baltic Sentry* announced January 2025, adding vessels, patrol aircraft and staff to Commander Task Force Baltic at Rostock.
- International Telecommunication Union and the International Cable Protection Committee – International Advisory Body for Submarine Cable Resilience (established late 2024); Abuja summit declaration on cable resilience (February 2025).
- United Nations Convention on the Law of the Sea, Articles 113-115, and the Convention for the Protection of Submarine Telegraph Cables (1884).

Not regulated financial advice.