# Ericsson Microwave Outlook

October 2026



ERICSSON

# Executive summary

## Contents

02 Executive summary
03 Backhaul capacity evolution
04 Microwave and satellite coexistence done right
08 Tackling the sunflower effect
12 Demystifying Layer 1 bonding

## Key contributors

**Executive editor:**
Git Sellin

**Articles:**
Andreas Olsson
Cristian Czegledi
Jonas Bjerger
Jonas Flodin
Jonas Hansryd
Magnus Nilsson
Mikael Coldrey
Mikael Öhberg
Pierre Hygren

5G growth and emerging 6G capabilities are increasing backhaul capacity demands across all radio access network (RAN) deployments. Capacity needs will vary widely by available spectrum, features, geography, and traffic patterns. Modern microwave solutions can efficiently support this evolution across both high-capacity and long-hop scenarios. In this year's report, we have updated the backhaul capacity table structure to cater for the introduction of 6G.

E-band microwave is critical for high-capacity 5G and future 6G backhaul, while satellite network operators increasingly seek to use the same spectrum for feeder links. However, with carefully defined satellite power limits, microwave links can be protected from satellite interference while still enabling satellite co-channel operations. With the World Radiocommunication Conference 2027 (WRC-27) set to decide these limits, our simulations show that coexistence is achievable through well-calibrated regulation and precise engineering.

Solar heating makes telecom towers bend away from the sun, creating the so-called sunflower effect. While the movement is small, it can significantly affect the E-band link performance. Tower tilt can be calculated or measured in the field, and its impact can be reduced through optimized antenna placement, tower design considerations, and different compensation technologies.

Multi-band scenarios benefit significantly from selecting the best aggregation method. This report compares hash-based schemes, policy-based forwarding, and Layer 1 bonding for aggregation of point-to-point microwave links in mobile backhaul networks. Layer 1 bonding delivers the most efficient solution, achieving more than 99 percent utilization across all tested bandwidth scenarios while preserving Quality of Service (QoS) and supporting large traffic flows.

# Backhaul capacity evolution

Current 5G and emerging 6G radio access network (RAN) deployments vary widely, creating a need for backhaul networks that can efficiently support RAN sites with diverse bandwidth requirements.

Mobile network data continues to increase, rising by 22 percent worldwide between Q1 2025 and Q1 2026, with 5G now carrying 50 percent of all mobile data traffic globally across more than 3 billion subscriptions[1]. Most of this growth is in 5G, as 5G population coverage, 5G mid-band coverage, and 5G mobile subscriptions increase worldwide[1]. Regional variations remain large with several markets below 50 percent as of 2025, indicating that the 5G expansion will continue well into the 2030s[1]. This will further increase the backhaul capacity needed at upgraded sites.

New services, such as artificial intelligence (AI) and smart glasses, are expected to drive a significant increase in uplink data compared with the current situation. Since most mobile backhaul networks offer the same downlink and uplink bandwidth, and uplink capacity is usually underutilized, this shift is not expected to have a significant impact on the transport networks or dimensioning guidelines.

6G is now on the horizon and the initial commercial services for early adopters are expected around 2030. While 6G offers many new network capabilities, the largest impact on backhaul network dimensioning will come from new spectrum assets in the centimetric range and higher modulation. Both will increase the achievable user data throughput.

Typical backhaul capacity requirements for distributed RAN three-sector end sites are shown in Figure 1 and calculated considering variations in deployment and usage.

Capacity variations remain significant, both between and within markets as well as across different deployment areas within service provider networks. This is driven both by the amount and type of spectrum available to service providers and their deployment choices, together with RAN features such as coordination services and multiple input, multiple output (MIMO) layers.

5G high-band deployments are limited to a few regions worldwide and are not included in the capacity table. High-band can significantly increase the required backhaul bandwidth. For example, three sectors of 800 MHz 4T/4R high-band can require more than 10 Gbps backhaul capacity, which must be considered if deployed.

The increased diversity in RAN deployments and features requires backhaul networks to offer efficient connectivity solutions across the full range of bandwidths shown in Figure 1. A thorough analysis of mobile network performance targets is key to finding the correct backhaul dimensioning and building the right transport network.

Modern microwave solutions can efficiently cover the wide range of bandwidths and geographical domains required in current and future mobile backhaul networks as described in Figure 1. They offer technologies such as E-band and multi-band for very high bandwidth needs, and wide channels and multi-carrier for longer hops with high performance. w

| | | 2026 | 2029 | 2032 |
|---|---|---|---|---|
| 4G & 5G | Urban | 1 Gbps - 3 Gbps | 3 Gbps - 10 Gbps | 5 Gbps - 15 Gbps |
| | Sub-Urban | 300 Mbps - 1 Gbps | 1 Gbps - 5 Gbps | 2 Gbps - 8 Gbps |
| | Rural | 150 Mbps - 500 Mbps | 300 Mbps - 2 Gbps | 500 Mbps - 3 Gbps |
| 4G, 5G & 6G introduction | Urban | 1 Gbps - 3 Gbps* | 3 Gbps - 10 Gbps* | 7 Gbps - 20 Gbps |
| | Sub-Urban | 300 Mbps - 1 Gbps* | 1 Gbps - 5 Gbps* | 3 Gbps - 10 Gbps |
| | Rural | 150 Mbps - 500 Mbps* | 300 Mbps - 2 Gbps* | 500 Mbps - 3 Gbps** |

* 6G spectrum introduction beyond 2029 ** No 6G spectrum additions expected in rural areas

Figure 1. Backhaul capacity per distributed site Source: Ericsson (2026)

[1] All subscription, traffic and commercial 5G SA launch information is from the Ericsson Mobility Report, June 2026

# Microwave and satellite coexistence done right

Satellite mega-constellations already use E-band, and more deployments are coming. WRC-27 will decide the power limits needed to keep microwave links protected. This article describes how that protection is achieved.

E-band backhaul continues to grow. Transceivers operating at 70/80 GHz nowadays account for a substantial percentage of the global installed base and are deployed in every region, according to previous editions of Ericsson Microwave Outlook. With 10 GHz of spectrum, E-band supports the multi-gigabit backhaul capacity needed for 5G and future 6G networks.

Satellite network operators want to use E-band spectrum for feeder links, the high-capacity connections between satellites and gateway ground stations. Satellite user terminals connect to satellites primarily in the Ku band (10.7–14.5 GHz) and Ka band (17.7–31.0 GHz) for broadband connectivity, while sub-3 GHz bands are primarily used for direct to mobile device connectivity. E-band is used only for connections between satellites and gateway ground stations. However, the scale is significant. Thousands of Non-Geostationary Orbit (NGSO) satellites, including Low Earth Orbit (LEO) and Medium Earth Orbit (MEO) satellites, already carry E-band radios, and gateway stations are operating in several markets. With up to 56,000 NGSO satellites planned across multiple constellations, alongside Geo-Stationary Orbit (GSO) satellites sharing the same spectrum, as many as 90 satellites may transmit simultaneously on the same frequency toward the same area where E-band backhaul links are deployed (see Figure 2).

About 90 satellites may transmit simultaneously on the same E-band frequency toward the same area.

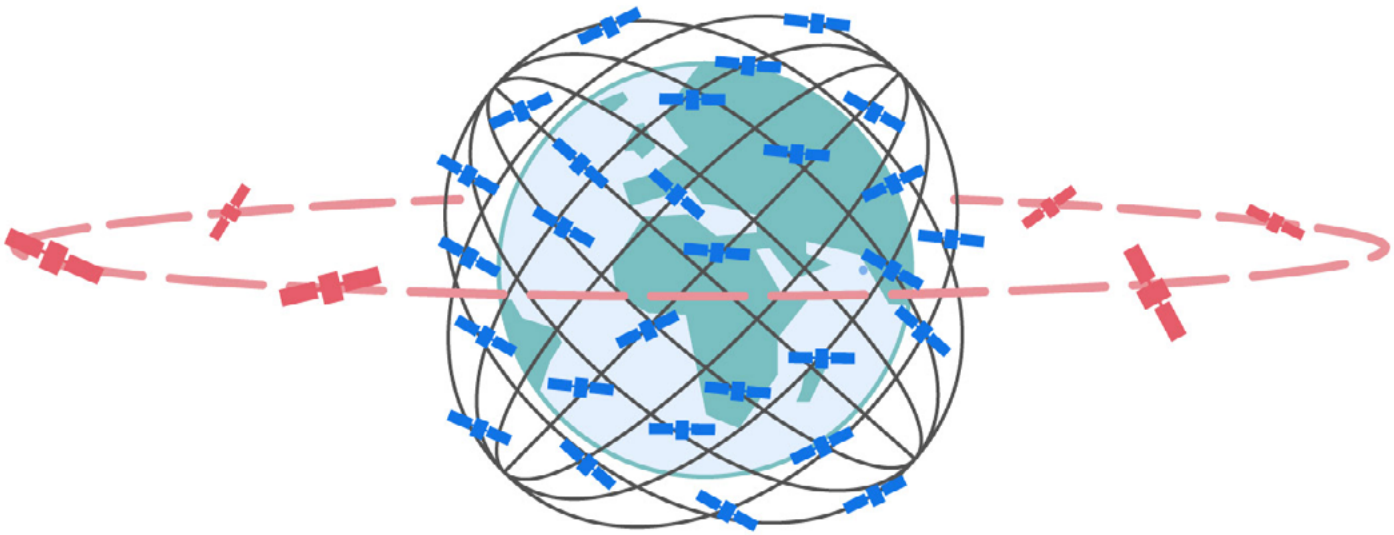


Figure 2. Global NGSO (blue) and GSO (red) satellite coverage with multiple orbital planes. Source: Ericsson (2026)

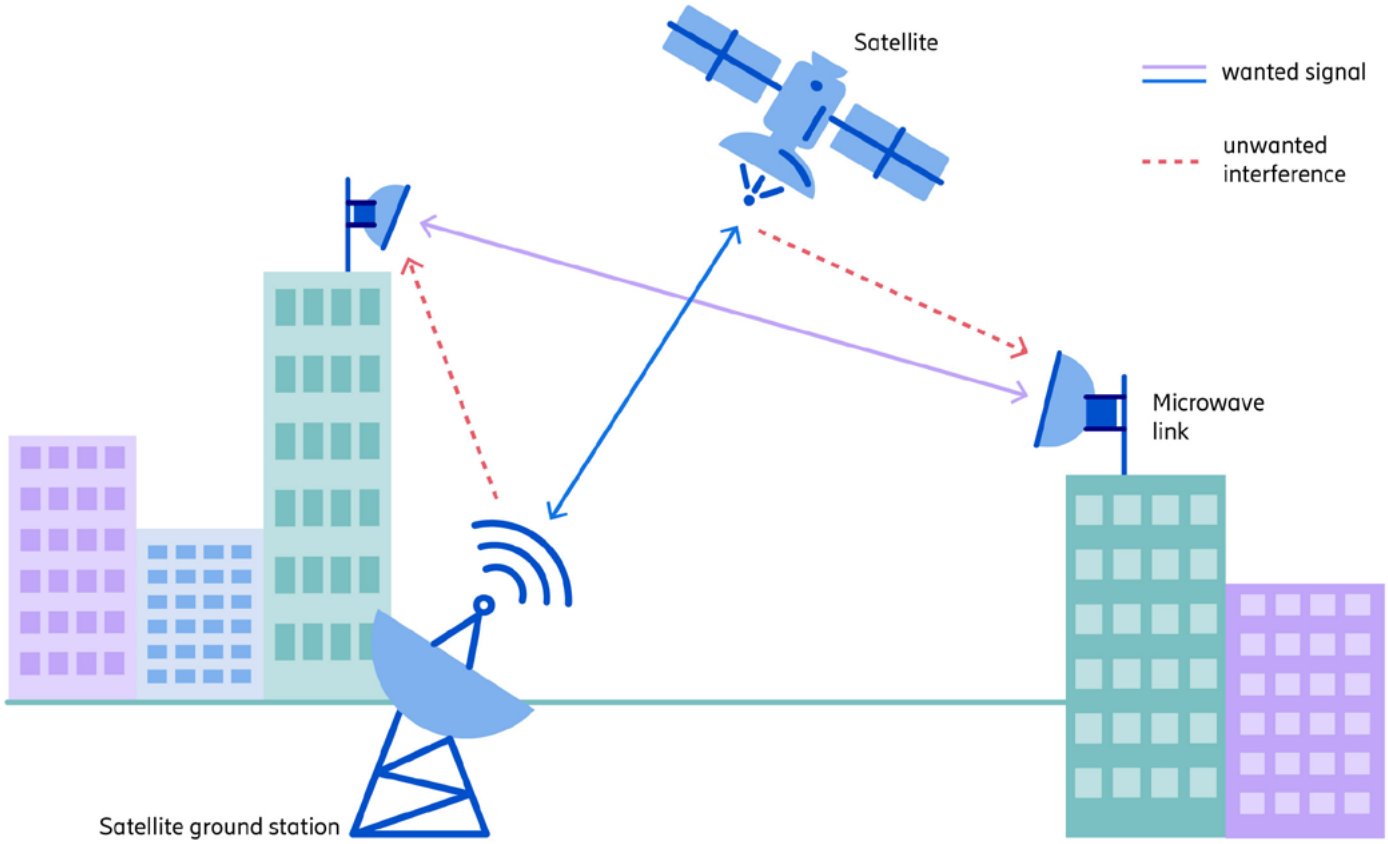



Figure 3. Interference geometry between satellite systems and a fixed service E-band link

Source: Ericsson (2026)

**The interference challenge**

Two satellite downlink paths are relevant, as shown in Figure 3. NGSO satellites transmit between 71 and 76 GHz. As they orbit, NGSO satellites periodically cross through the narrow main beam of a Fixed Service (FS) microwave antenna. Each crossing lasts a few seconds and can produce an intense interference spike. GSO satellites remain fixed on the geostationary arc and produce constant interference to microwave links, with the highest levels affecting links pointing toward the arc. A third path comes from satellite ground stations transmitting in 81-86 GHz, radiating high power that couples into nearby microwave receivers.

Geometry drives everything. A microwave antenna at 0 degrees elevation has additional protection against satellite downlink interference than one at 10 degrees. Two effects combine: the antenna pattern rolls off steeply away from boresight, and the longer atmospheric path at low angles attenuates the satellite signal. Most E-band links point near the horizon and benefit from both effects.

**A manageable challenge**

Spectrum sharing is not new to microwave backhaul. Wi-Fi and NR Unlicensed have moved into parts of the 6 GHz backhaul band, raising real questions about interference to licensed links. These devices have power limits, and some countries require automated frequency coordination for standard-power operation. However, they operate without individual frequency assignments or known deployment locations. Managing aggregate interference from many uncoordinated transmitters can be challenging.

The satellite case is different. Here, the interference comes from a licensed system, not from an open population of consumer devices. The satellite network operator declares its constellation in advance, including the number of satellites, their power levels, and antenna characteristics, while staying individually accountable for it. The set of emitters is bounded and known, so the aggregate energy reaching the ground can be computed and capped in advance, rather than managed after the fact. That is what a Power Flux Density (PFD) mask does: it limits how much interference a satellite may deliver to the Earth's surface, protecting microwave links while still enabling satellite operations. This approach has precedent. PFD masks have protected fixed links from satellite downlink interference for decades in many bands, including Ku and Ka, as established in Article 21 of the ITU Radio Regulations. At E-band, the same PFD approach applies to satellite downlinks in 71-76 GHz. What is new about 71-76 GHz is the band, not the method. The question is where to set the numbers.

Satellite ground stations present a separate challenge. They can transmit at high power, and their interference propagates near the horizon, exactly where microwave antennas point. Satellite uplink is regulated by existing Equivalent Isotropically Radiated Power (EIRP) limits and studies show that a 27-50 km separation is needed between satellite ground stations and microwave receivers in worst-case alignment.

**Regulatory studies and protection criteria**

Scheduled for November 2027, the World Radiocommunication Conference (WRC-27) will establish PFD limits for satellite downlinks in E-band. The PFD limits will be incorporated into Article 21 of the Radio Regulations. These limits will cover all satellite services sharing these bands (fixed-satellite, mobile-satellite, and broadcasting-satellite) as mandated by Resolution 775 (Rev.WRC-23). In addition, Appendix 7 of the Radio Regulations may need to be modified to facilitate international coordination with satellite ground stations considering the new band. Studies are being carried out in ITU-R Working Party 5C, with contributions from administrations including Europe, China, Canada, India, the Russian Federation, and the United States. The satellite and mobile communications industries are also actively engaged. The draft Conference Preparatory Meeting (CPM) text was finalized in May 2026.

Two protection criteria have been agreed:

- **Long-term:** The aggregate interference-to-noise ratio (I/N) at the microwave receiver must stay below -10 dB for more than 20 percent of the time.
- **Short-term:** I/N must not exceed +11 dB for more than 0.00128 percent of the time (about one second per day).

The long-term criterion limits how much satellite interference can eat into a link's fade margin day after day. The short-term criterion caps the brief peaks that may cause momentary outages. If PFD limits are set too loose, operators face reduced margins on exposed links, lower availability, or shorter hops.

The proposed PFD limits depend on elevation angle: they are strict at low angles, where microwave antennas point, and more relaxed at high angles, where antenna gain toward the satellite drops off.

**What our simulations show**

Both the short- and long-term protection criteria must be met simultaneously. The nature of each satellite type, however, determines which one is reached first. There are multiple proposals for PFD masks being discussed, ranging from strict to relaxed, depending on system assumptions. The PFD limit considered in this study is one proposed by the FS industry.

GSO satellites are stationary. They produce constant interference toward all microwave links in their footprint, with the highest levels affecting links whose beams point toward the geostationary arc. No variation means no short-term peaks, so meeting the long-term criterion automatically satisfies the short-term one. Our simulations with satellites at 1-degree spacing show that aggregate GSO I/N stays well below -10 dB at most orientations (Figure 4). The worst cases are specific azimuths where the microwave antenna beam points straight at the arc. Larger antennas have higher gain toward the arc and therefore collect more GSO interference. The simulation uses a 0.6 m antenna, a very common size in E-band deployments, making the 0.6 m antenna the critical case for GSO protection.

NGSO is different. Satellites move constantly, producing random interference and brief spikes each time one crosses the microwave main beam, which is captured by the short-term protection criterion. Our simulations confirm that the considered PFD mask satisfies both the long-term and the short-term criteria, although the margin is tight at certain azimuth orientations, as shown in Figure 5 where the dashed red region represents the protection criteria.

Figure 5 shows the scenario when the FS receiver is placed on the equator and pointing in different azimuth directions, with satellites distributed across all orbital planes. The worst case occurs when the microwave antenna is on the equator pointing either east or west: every satellite orbit crosses the equator regardless of inclination, so the equatorial receiver is exposed to satellites from all orbital shells simultaneously. Moving the receiver toward higher latitudes reduces the number of visible shells and increases the margin.

The simulations use a 0.2 m antenna. Antenna size determines which protection criterion is more stringent. Smaller antennas have higher sidelobes that collect more aggregate interference from all visible satellites, pushing toward the long-term limit. Larger antennas shift the constraint toward the short-term criterion: their higher main beam gain amplifies each beam-crossing spike, particularly when the antenna is pointed toward an orbit where crossings are most frequent.

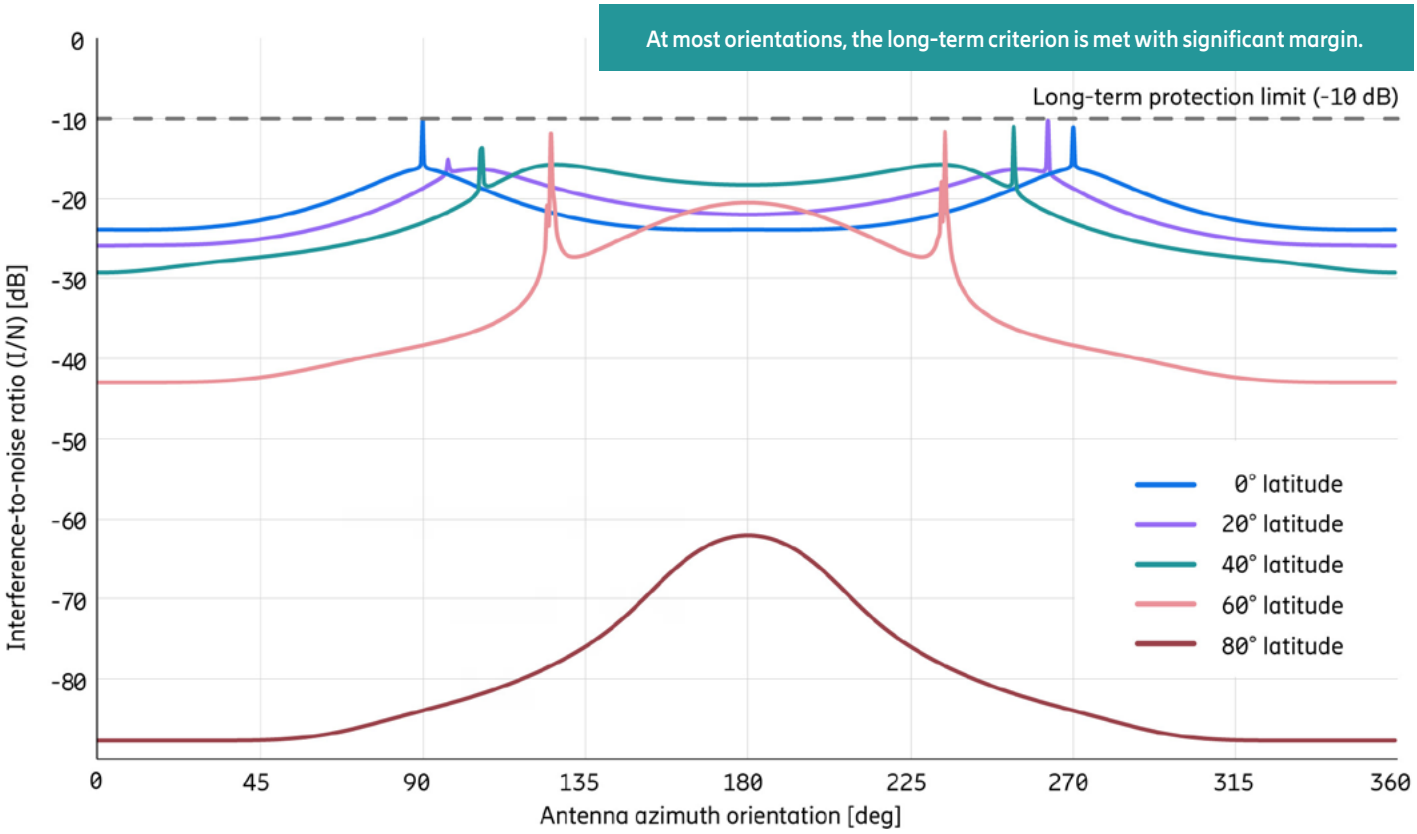



Figure 4. Simulated I/N from aggregate GSO interference at different FS locations and azimuth orientations

Source: Ericsson (2026)

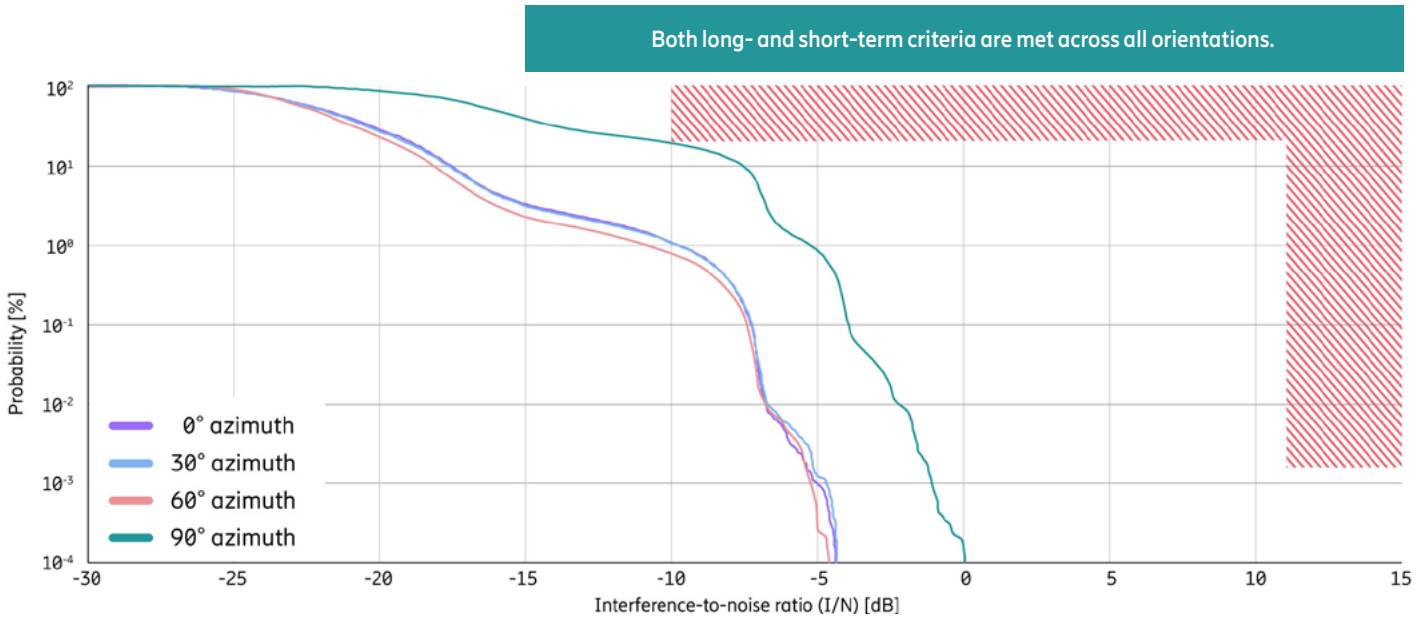



Figure 5. Simulated complementary cumulative probability distribution of I/N from NGSO systems at different azimuth angles Source: Ericsson (2026)

| Simulation assumption | | |
|---|---|---|
| Parameter | GSO (Figure 4) | NGSO (Figure 5) |
| FS antenna | 0.6 m / 51 dBi | 0.2 m / 41.5 dBi |
| FS elevation angle | 10° | 10° |
| Frequency | 75.75 GHz | 75.75 GHz |
| Antenna pattern | ITU-R F.1245 | ITU-R F.1245 |
| Satellite systems | GSO at 1° spacing, full arc (360 satellites) | NGSO with 54 shells (~19700 satellites) |
| Atmospheric attenuation | Included | Included |

**Conclusion**

The high atmospheric loss at E-band for low-elevation satellite communications makes this type of operation unlikely. This opens a path forward: stricter PFD at low-elevation angles need not constrain satellite communications in practice, and our simulations show that PFD limits can be set to satisfy the FS protection criteria.

While the final values for the PFD mask remain to be agreed at WRC-27, the evidence from ITU-R studies is clear: coexistence between satellite mega-constellations and E-band microwave is possible. The satellite scenario has known parameters and predictable geometry, making it suitable for precise engineering. With a well-calibrated PFD mask at WRC-27, Communication Service Providers (CSP) can invest in high-capacity E-band backhaul with confidence that the spectrum will stay protected.

Ultimately, the two services are complementary. E-band microwave delivers carrier-grade, multi-gigabit transport in urban and suburban areas for 5G and future 6G networks. Satellite feeder links at E-band support broadband services reaching areas beyond terrestrial network coverage. Well-designed regulatory limits let both thrive in the same spectrum, each serving the areas where it is most needed.

# Tackling the sunflower effect

A small temperature difference across a tower can create challenges for narrow beam microwave links. Learn why the sunflower effect matters and how to tackle it.

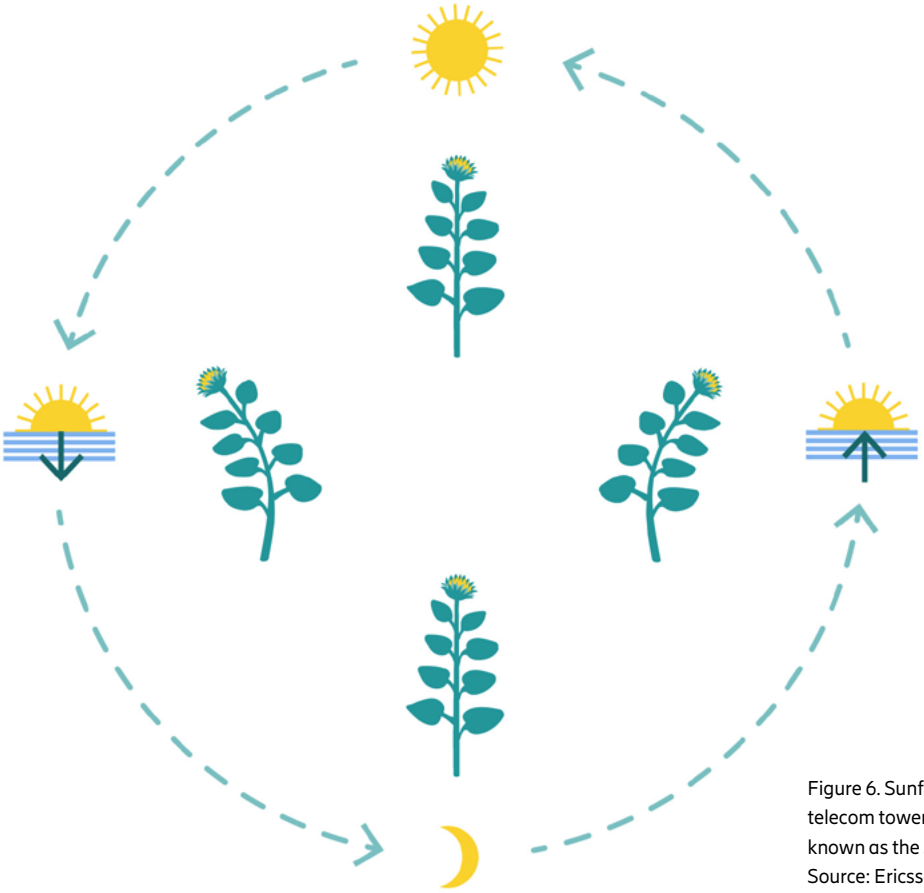


Figure 6. Sunflowers bend toward the sun while telecom towers bend away from the sun. This is known as the sunflower effect.
Source: Ericsson (2026)

For Communication Service Providers (CSPs) deploying E-band links, small tower movements can have a significant impact. A narrow beam that shifts by fractions of a degree can lose signal strength, reduce fade margin, and interrupt services. One cause is the sunflower effect. Understanding and managing this effect is particularly important for high-capacity microwave links that depend on precise antenna alignment.

The sunflower effect starts when sunlight heats one side of a tower. The warmer side expands, the tower bends away from the sun, and the antenna points slightly off target. For narrow E-band beams, that small movement can be enough to affect performance.

Towers are designed to remain stable under heavy loads and harsh weather conditions, but no material is immune to temperature changes. When one side expands more than the other, alignment can shift. Table 1 shows thermal expansion values for different materials.

Uniform heating rarely creates problems. The risk arises when one side of the tower is warmer than the other, especially in direct sunlight. The hot side expands more than the cold side, creating strain and a slight tilt.

| Material | $\alpha$ [$10^{-6}$/°C] |
|---|---|
| Cast iron | 10.4-11 |
| Steel | 10.8-12.5 |
| Concrete | 13-14 |
| Aluminum | 21-24 |

Table 1. Thermal expansion constant, $\alpha$ for different materials.

**Estimate how much the tower moves**
To understand when the sunflower effect becomes operationally significant, it helps to estimate how much tower movement a given temperature difference can cause.

To estimate this, the tower may be divided into small segments, each with height **x**, When the sun warms one side of a segment, each segment expands by **Δx** (see Figure 7), calculated as:

$$\Delta x = x\alpha\Delta T$$

Equation 1

Here, **ΔT** is the temperature difference between the two sides of the segment. For each segment the expansion on the warmer side causes an angular tilt, **φ**.

If **D** is the tower diameter, **φ** can be calculated as:

$$\varphi = \tan^{-1}\frac{\Delta x}{D}$$

Equation 2

A tower with height **h**, may be divided into **n** small segments, each with height **x**. The total tower tilt **θ** at the top is the sum of the angular tilt **φ** across all segments. As the segments get smaller and **n** approaches infinity, the aggregate tilt in radians can be calculated using Equation 3, giving the aggregate tower tilt **θ** in radians. Figure 7 shows how the tower will bend away from the sun in a curve.

$$\theta = \frac{h\alpha\Delta T}{D}$$

Equation 3

Using Equation 3 we can see that a 35 m steel monopole with a 0.6 m diameter, and a 10°C temperature difference can create a top-of-tower tilt of 0.40°. In many E-band deployments, that is enough to misalign the antenna and degrade performance.

Example 1

**Measure tower movement without climbing**
Tilts due to sunflower effect are small and can fall below the resolution of standard gyroscope tools, while climbing the tower to install measurement equipment is often impractical and costly. A commercial laser distance meter gives field teams a simpler option: measuring and logging horizontal movement at the top of the tower from the ground.

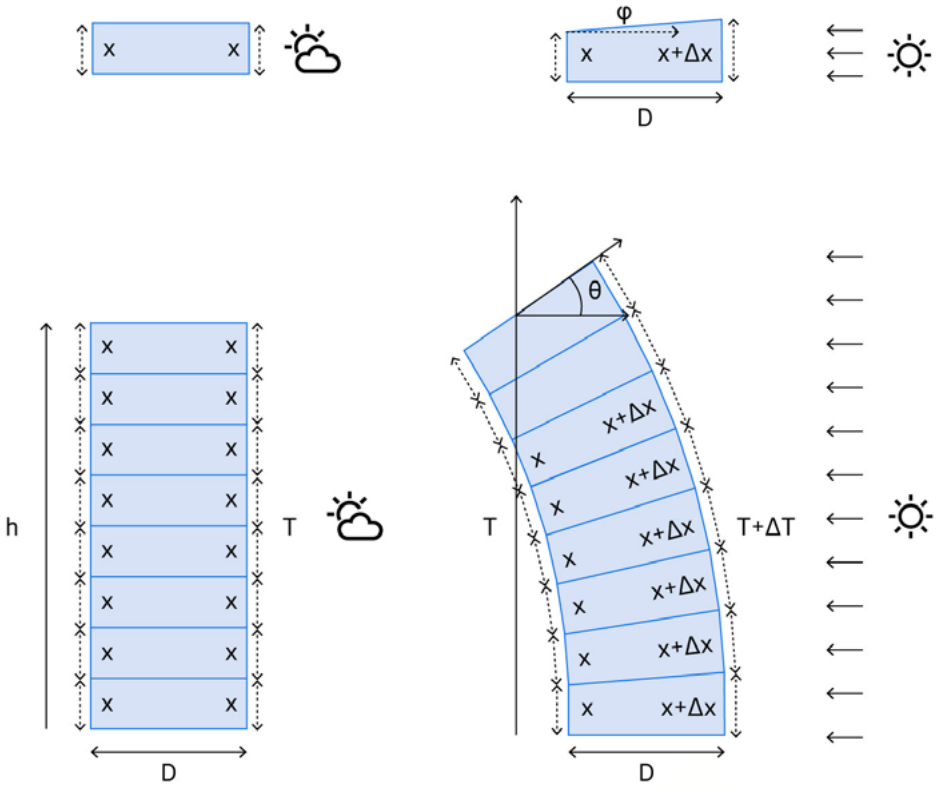



Figure 7. Aggregated tilt, **θ** of the tower when exposed to heat. A material exposed to heat will expand in length.
Source: Ericsson (2026)

Using a distance meter, we measure a deviation **ΔZ** from the centerline. Using the same model as above, based on segmentation we can estimate the deviation as the sum of movement across all segments. Equation 4 gives the closed-form expression.

$$\Delta Z = \frac{a\Delta T h^2}{2D}$$

Equation 4

Combining Equation 3 and Equation 4 give the tower tilt **θ** in radians as a function of the measured tower deviation **ΔZ**.

$$\theta = \frac{2\Delta Z}{h}$$

Equation 5

To assess the sunflower effects on installed telecom equipment in a 35 m high steel tower, a service engineer places a laser distance meter on the ground near the site and records top-of-tower movement for 24 hours. The next day, the logs show a maximum movement of 13.8 cm. Using Equation 5, this corresponds to a maximum tilt of 0.45°.

Example 2

A common mistake is to treat tower movement as a straight-line tilt from the base to top. The tower bends in a curve, as shown in Figure 7. The straight-line approximation underestimates the actual tilt by a factor of two.

CSPs do not always need dedicated measurement equipment to identify sunflower effects. Long-term performance analytics from the live microwave link can reveal recurring alignment changes and help estimate their impact on network performance. Figure 8 (bottom) shows attenuation measured over one week on an E-band link installed on a 22 m steel monopole. The daily pattern, shown in detail in the top panel, is consistent with solar heating of the tower. As one side of the pole warms, the antenna moves off alignment, and attenuation increases. Comparing the measured attenuation with the antenna pattern indicates a maximum tilt of approximately 0.8°. For a monopole with a diameter of 30 cm, this corresponds to a temperature difference between front and back of about 16°C.

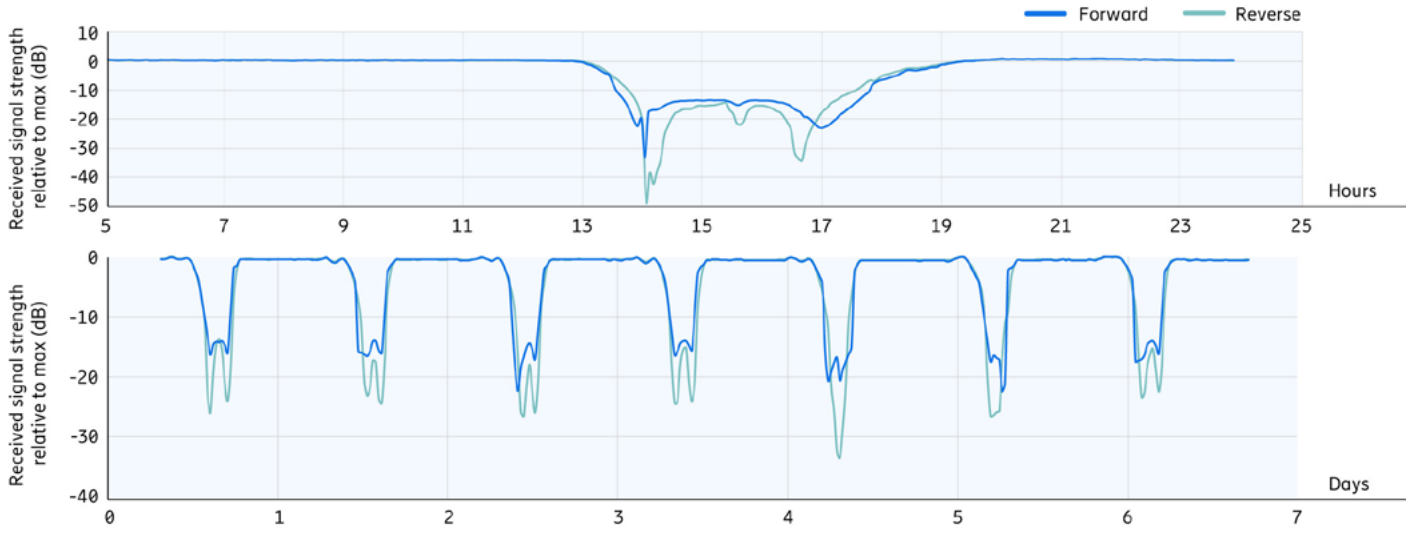



Figure 8. Measured long-term attenuation over a week for an E-band monopole installation with 0.6 m antenna (bottom). Measured attenuation over 24 hours (top). Source: Ericsson (2026)

**Why sunflower effects matter for E-band links**

E-band links deliver very high capacity using narrow, highly directional beams. That makes them particularly sensitive to tower movement. Even small changes in antenna alignment can reduce received power, shrink fade margin, and affect service availability.

A 0.6m E-band antenna can deliver more than 50 dBi of gain, but its half-power beamwidth is only 0.38°. A tilt of 0.19° can cut received power by 3 dB, and a 0.35° misalignment can reduce antenna gain by more than 10 dB. With 0.6 m and larger E-band antennas, tower stability becomes part of the link budget, not a site detail.

A tilt of 0.19° can reduce received power by 3 dB, and a 0.35° misalignment can reduce antenna gain by

~10dB

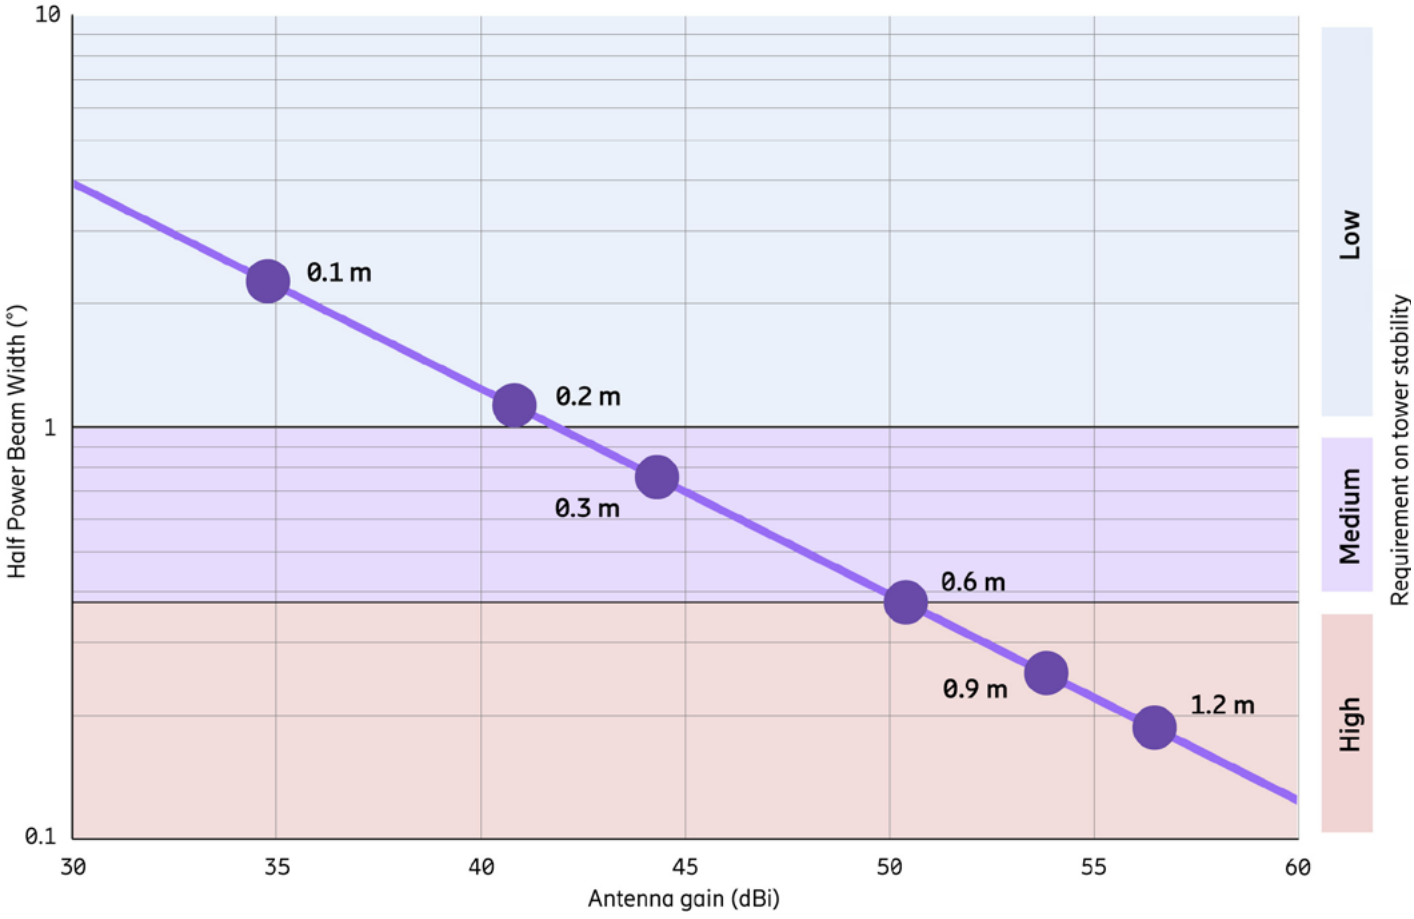



Figure 9. Half Power Beam Width versus antenna gain for different diameters of E-band reflector antennas. Source: Ericsson (2026)

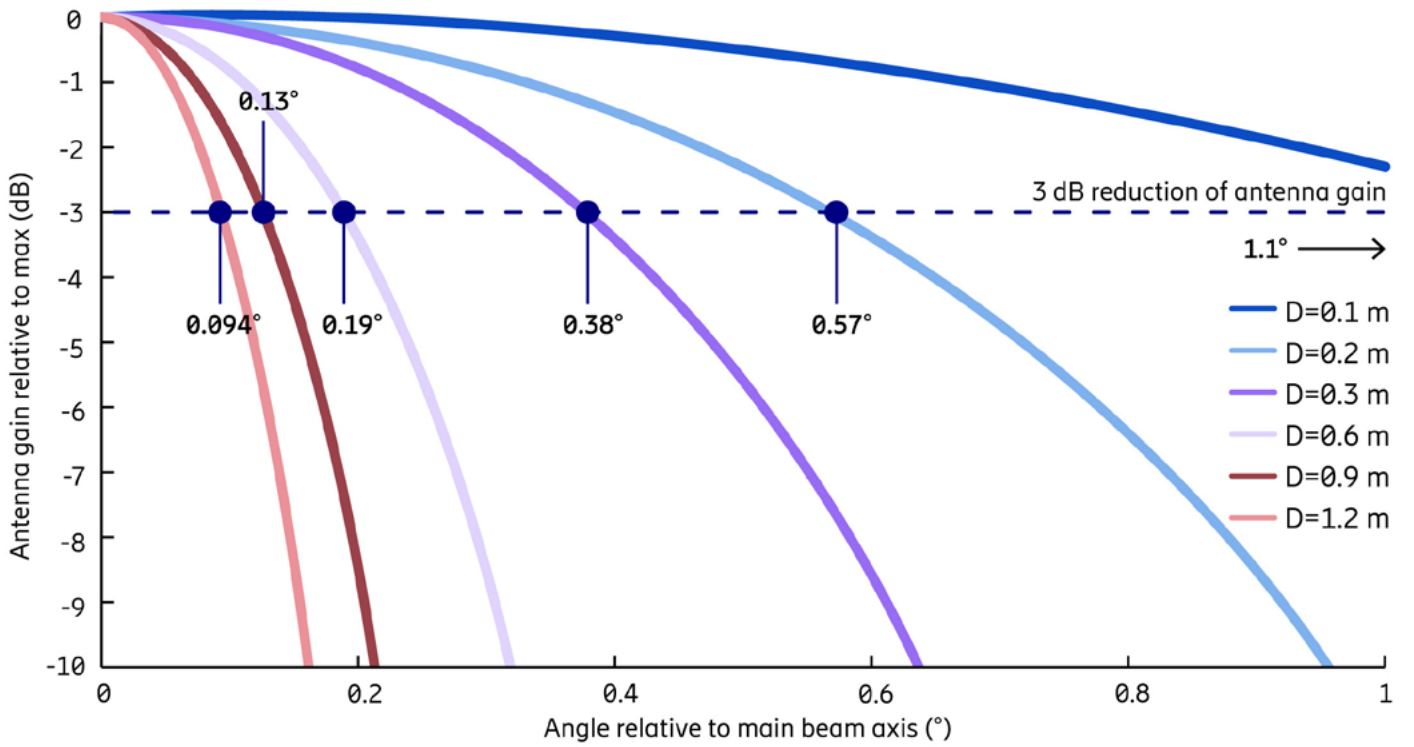



Figure 10. Relative E-band antenna gain versus misalignment for 0.1 m, 0.2 m, 0.3 m, 0.6 m, 0.9 m, and 1.2 m diameter antennas. Source: Ericsson (2026)

Tower bending caused by the sunflower effect is proportional to the temperature difference between the hot and cold sides of the tower. If the temperature difference increases from 10°C to 20°C, the misalignment doubles from 0.44° to 0.88°. This can sharply reduce antenna gain and, in some cases, cause a link outage.

**Reduce the impact on live links**

CSPs have two main choices: reduce how much the tower moves, or keep the antenna aligned as the tower moves. The right approach depends on the site, the link budget, and how much variation the service can tolerate.

Equation 3 shows that angular tilt is proportional to installation height. Placing the antenna halfway down the tower cuts the tilt by about 50 percent. A larger tower diameter also reduces tilt. Because the sunflower effect depends on the temperature difference across the structure, it is less common in lattice towers and more common in monopole installations.

When site conditions cannot be changed, the radio link can compensate for tower movement. One option is a sway-compensation antenna that can correct mast movement of ±2° or more. A common implementation steers the feeder inside the antenna. This supports both sunflower effect tilt and wind-induced sway. The trade-off is that antenna gain can still vary because efficiency decreases as the feeder angle changes.

Another option is an electrically steerable antenna fixture that continuously realigns the antenna as tower tilt changes. This can support larger angular deviations, and one fixture can match different radio and antenna combinations. The trade-offs are response speed and mechanical exposure. The combined radio and antenna mass makes fast wind-sway correction difficult, and the mechanical installation can be more vulnerable to wind, snow, and ice.

As microwave transport carries more capacity over narrower beams, tower movement becomes a network performance issue. By combining field insight, link analytics, and alignment compensation, CSPs can protect fading margins and keep services available. The microwave industry continues to develop microwave transport solutions that help CSPs maintain alignment, protect link performance, and maximize network availability under all types of challenging environmental conditions.

Placing the antenna halfway down the tower cuts the tilt by about

**50%**

# Demystifying Layer 1 bonding

Not all link aggregation methods deliver the capacity they promise. Traditional aggregation methods can leave valuable microwave spectrum underutilized.

For communication links, aggregating two or more links to increase data transmission capacity is well known. Several methods exist, and this article describes and compares three different alternatives for aggregating point-to-point microwave links in backhaul networks:

- Hash-based distribution schemes such as Link Aggregation (LAG) and Equal Cost Multi Path (ECMP)
- Policy Based Forwarding (PBF)
- Layer 1 (L1) bonding, also commonly known as Radio Link Bonding (RLB)

There are different scenarios in which data is combined over multiple carriers in microwave networks. This evaluation uses a multi-band scenario, which is increasingly being deployed and benefits significantly from using the best aggregation method. One 15 GHz standard band link is combined with an E-band link to form a high-capacity multi-band group (see Figure 11).

Multi-band scenarios benefit significantly from using the best aggregation method.

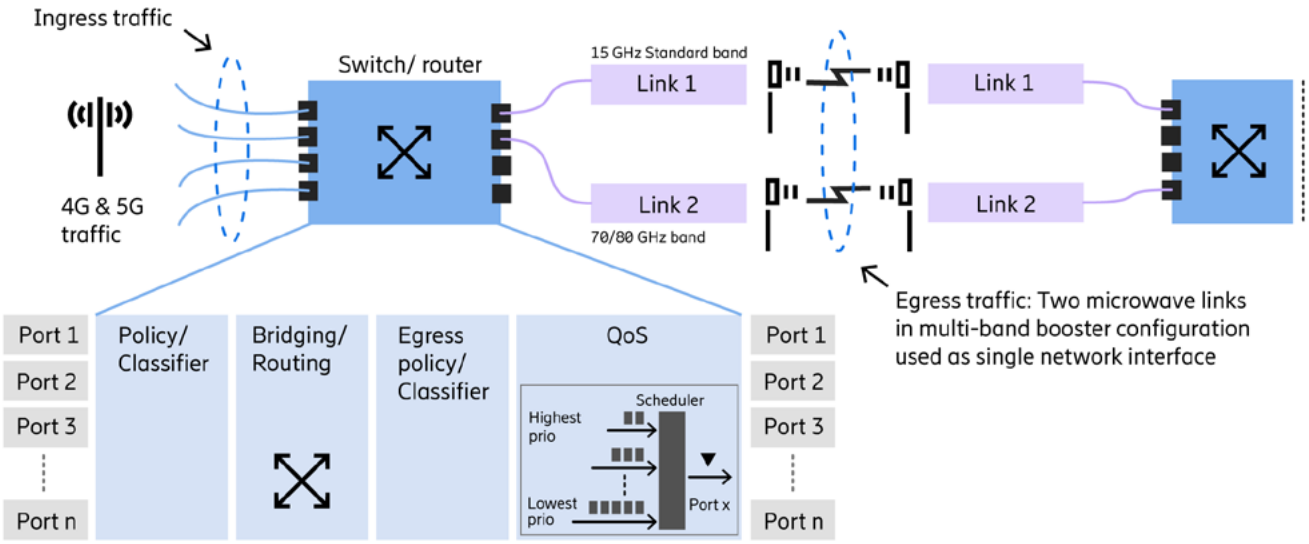



Figure 11. Aggregation test set-up

Source: Ericsson (2026)

**Hash based schemes**

LAG and ECMP are hash-based schemes in which link or route allocation is performed based on flow identification derived from higher-layer protocol address fields. Layer 2/Layer 3 (L2/L3) hash-based schemes are less complicated to implement and are often natively supported in switch and router hardware. In practice, a flow is identified by its hash checksum value, typically calculated from static address fields, and then assigned to a physical link in an aggregation group. Subsequent data with the same hash checksum value is then forwarded to the originally assigned link (see Figure 12).

This method has several limitations:

- Statistical bias: The hashing algorithm can interfere with address assignment rules in the network, causing a systematic preference for one link.
- Few flows: If the number of flows is not enough (very common problem in backhaul, especially when traffic is encrypted) overall utilization can be low even if hashing is optimal.
- QoS impact on individual congested links: Data can be dropped even if there is capacity available, and even worse, high-priority data can be dropped over low-priority data.
- Flow size limitation: An individual flow can never be larger than the capacity of a specific link in the aggregation group. This is a significant limitation in backhaul, where few and large flows are common.

There are different ways to overcome some of the limitations described. Typical mitigations include:

- Dynamic load balancing: A function that continuously measures the utilization of each link in the aggregation group and rehashes to optimize utilization. Silent periods are used to avoid reordering at the receiving end when flows are moved between links.
- Adding weight as an input to the hash calculation: This can improve utilization in an aggregation group where the links have different steady-state bandwidths.

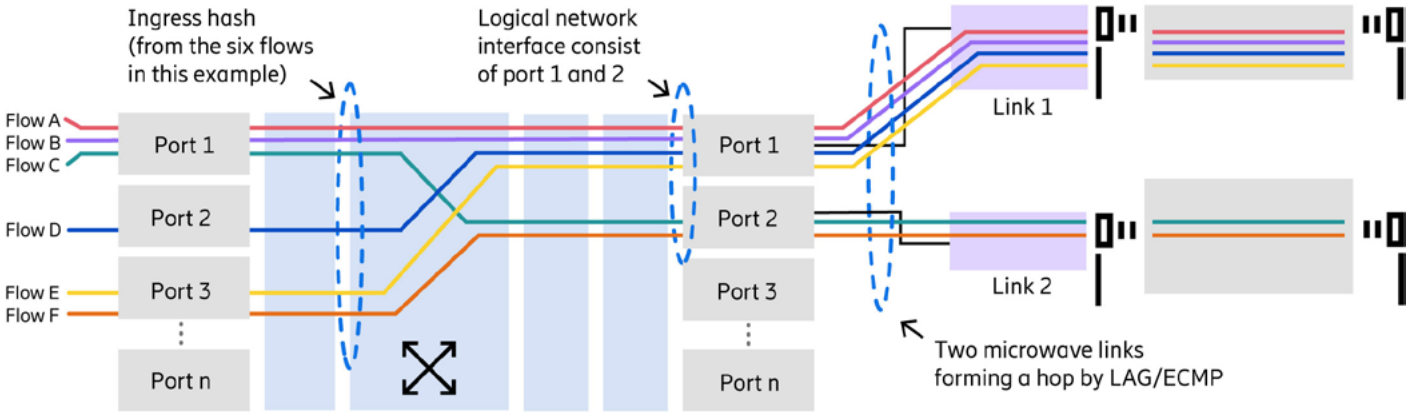



Figure 12. Link Aggregation (LAG) or Equal Cost Multi Path (ECMP) flows

Source: Ericsson (2026)

**Policy Based Forwarding**

PBF means that flows are assigned to a link based on static address fields, as in L2/L3 hashing methods. The difference is that PBF explicitly assigns flows based on predefined policies aligned with service priorities. A typical example would be a multi-band aggregation group where one standard band link and one E-band link form an aggregation group, and the high-priority traffic is protected instead of optimizing the utilization. This means that high-priority traffic is assigned to the link with high availability, while the rest of the traffic is assigned to the high-capacity link, which has lower availability. See Figure 13 for the comparison set-up.

PBF can also be combined with dynamic control, in which flows are moved between links in a controlled manner and utilization is measured after each adjustment.

PBF shares some of the limitations inherent in hash-based schemes:

- Few flows: If the number of flows is low, the overall utilization can be low.
- Flow size limitation: An individual flow can never be larger than the capacity of a specific link in the aggregation group.

One advantage of PBF over hash-based schemes is that it can be configured to provide much better protection for high-priority traffic.

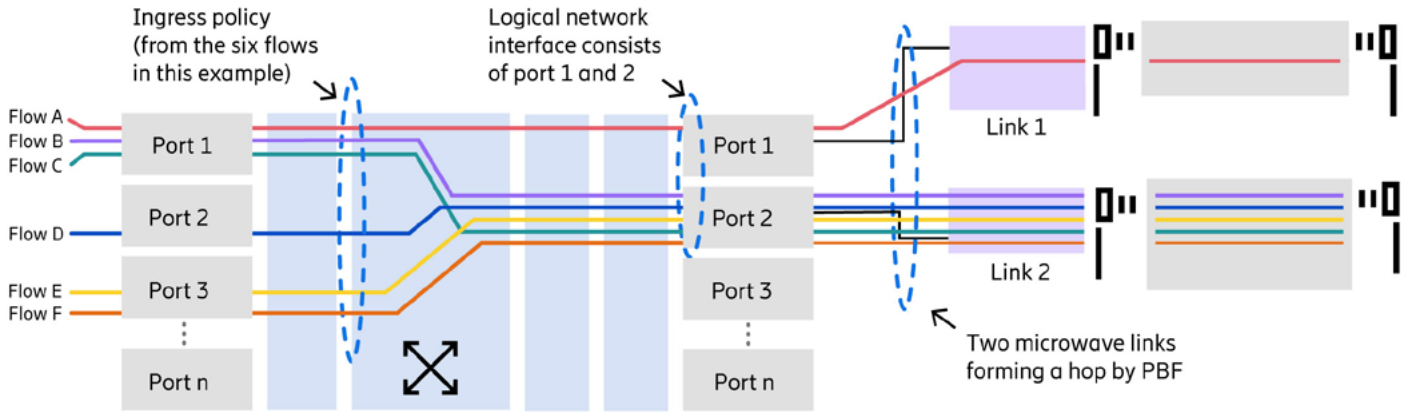



Figure 13. Policy based forwarding (PBF) flows

Source: Ericsson (2026)

**Layer 1 bonding**

L1 bonding or RLB, is fundamentally about spectrum efficiency. L1 bonding operates below the data link layer and is therefore transparent to Layer 2 and higher protocol layers. Any change to the network architecture, such as adding encryption, will not affect the efficiency of L1 bonding.

Bonding means that different parts of the traffic are transmitted over different links and arranged in order when received.

The core mechanism in L1 bonding is that the transmitted traffic is distributed across the links. The distribution function also applies a sequence number to each data frame: this is needed to keep track of the frames on the receiving side. At the far end, the receiving function allows for different arrival times between the links. It then puts the frames together in the correct order to avoid out-of-sequence delivery of frames in the network (see Figure 14).

L1 bonding enables the service providers to aggregate two or more carriers while managing them as a single network interface from a Layer 2, Layer 3, and QoS perspective. The aggregated capacity for the group is always perceived and managed as a single resource:

- The overall bandwidth is the sum of the included carriers.
- QoS ensures that high-priority traffic is not dropped over lower-priority traffic.
- A single flow is never limited to the capacity of a specific link. If the flow has a bandwidth greater than the capacity of one link in the group, the data is transmitted over several links.

L1 bonding is designed to maximize the bandwidth a service provider can extract from the available spectrum, without the need to consider or plan for traffic diversity in the network.

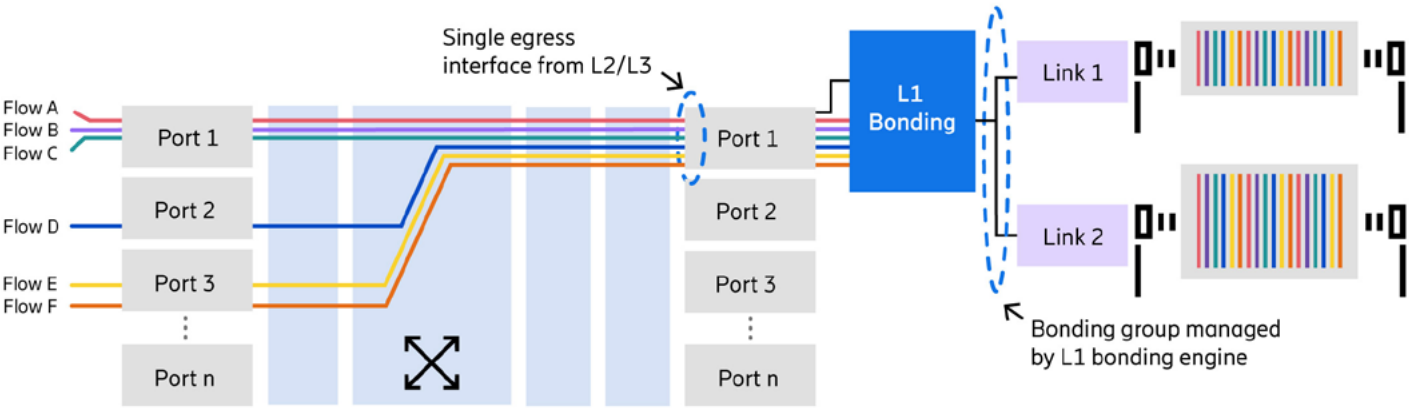



Figure 14. Layer 1 Bonding flows

Source: Ericsson (2026)

**Putting aggregation methods to the test**

To illustrate and compare the methods described above, the following assumptions are used. The scenario is a co-located 4G/5G site, where both a long-term evolution (LTE) eNodeB and a new radio (NR) gNodeB are physically installed at the same tower or rooftop location and use the same backhaul connection to the core network, IPsec encrypted. The 4G layer provides wide-area coverage and a mobility anchor, while 5G provides high throughput in denser coverage areas.

The coverage radius is typically a few hundred meters to a few kilometers, depending on the frequency band and environment. The combined backhaul demand is approximately 5–8 Gbps at peak (typically three sectors, each with both 4G and 5G antennas).

For this evaluation, the site's network traffic is categorized as follows:

| Priority class | PCP | Traffic content | Percentage of total volume | Visible flow count |
|---|---|---|---|---|
| Highest | 7 | PTP timing | <0.1% | 1 |
| High | 5-6 | VoLTE/VoNR, S1-MME/N2, X2/Xn signaling | 5-10% | 2-4 |
| Medium | 2-3 | OAM, management | <0.5% | 1-2 |
| Low | 0-2 | Bulk user data (best effort) | 85-95% | 1-2 |

The microwave configuration uses a multi-band setup: a 15 GHz standard-band link with 56 MHz channel spacing in a 2+0 XPIC configuration, providing around 1.1 Gbps at the highest modulation, combined with a 70/80 GHz E-band link with 750 MHz channel spacing in a 1+0 configuration, providing around 5.1 Gbps. Together, the hop provides an aggregated bandwidth of approximately 6.2 Gbps.

During the test, different fading scenarios are configured, and overall utilization and packet drops (controlled and uncontrolled) are measured and compared, as shown in Figure 15.

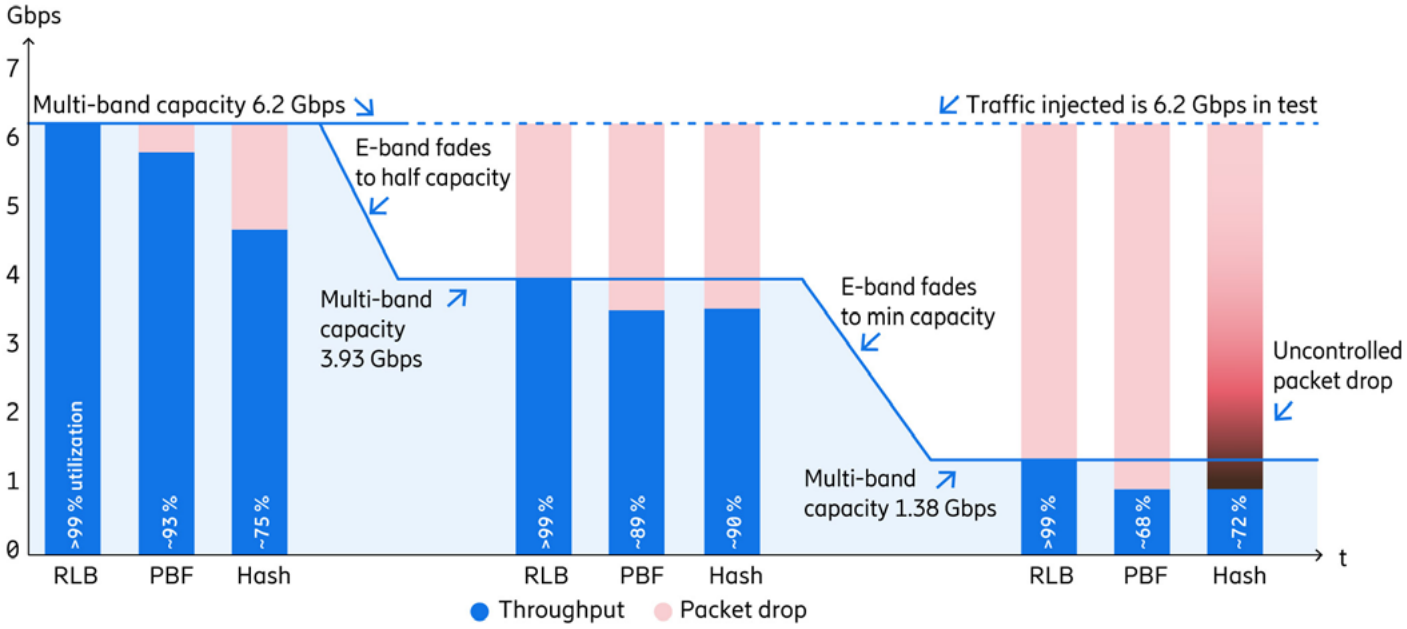



Figure 15. Throughput comparison during different fading events

Source: Ericsson (2026)

**Turning scarce spectrum into usable capacity**

The evaluation shows that hash-based schemes can be useful under specific conditions, but there is a correlation between backhaul traffic diversity and what the overall utilization a group can reach. As shown in the test results, utilization varies between 72-90 percent for hash-based schemes. When the E-band link drops to a very low bandwidth, uncontrolled packet drops also begin to occur. This is because the flows do not map well to the available bandwidth, and high-priority traffic cannot be protected. To achieve good utilization when using hash-based schemes, the number of flows needs to be sufficiently large, and the bandwidth of each flow should be small compared with the bandwidth of the links. While this condition is met in some networks, it is more difficult to achieve in mobile backhaul environments. Combined with the risk of suboptimal QoS handling and the limitation that each flow is restricted to the capacity of a single link, hash-based schemes are not well suited for point-to-point microwave links where spectrum is a scarce resource.

PBF utilization depends heavily on how well each flow aligns with the bandwidth of its assigned link within the PBF group. In this test scenario, the flows align relatively well, allowing PBF to reach approximately 90 percent utilization. When the E-band link fades to very low bandwidth, utilization drops to around 68 percent. PBF offers a slight improvement over hash-based schemes by protecting high-priority traffic, as the defined policy is well suited to a multi-band setup. However, in aggregation groups with links of similar bandwidth, such as standard-band 2+0 or 4+0 configurations, PBF is not suitable. This is because steering high-priority traffic to a specific link offers no real advantage in an aggregation group where the links have similar bandwidth and availability. As a result, performance may even be worse than with hash-based schemes.

L1 bonding achieves more than 99% utilization across all bandwidth scenarios.

L1 bonding delivers superior spectrum efficiency and traffic handling across all microwave backhaul scenarios, regardless of link speed differences or traffic diversity. Compared with hash-based and policy-based methods, it maximizes link utilization while preserving QoS and supporting large flows. As demonstrated in this test, it achieves more than 99 percent utilization across all bandwidth scenarios.

ericsson.com

The content of this document is subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

EAB-26:006958 Uen
© Ericsson 2026