Ericsson Microwave Outlook

October 2026



ERICSSON

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.

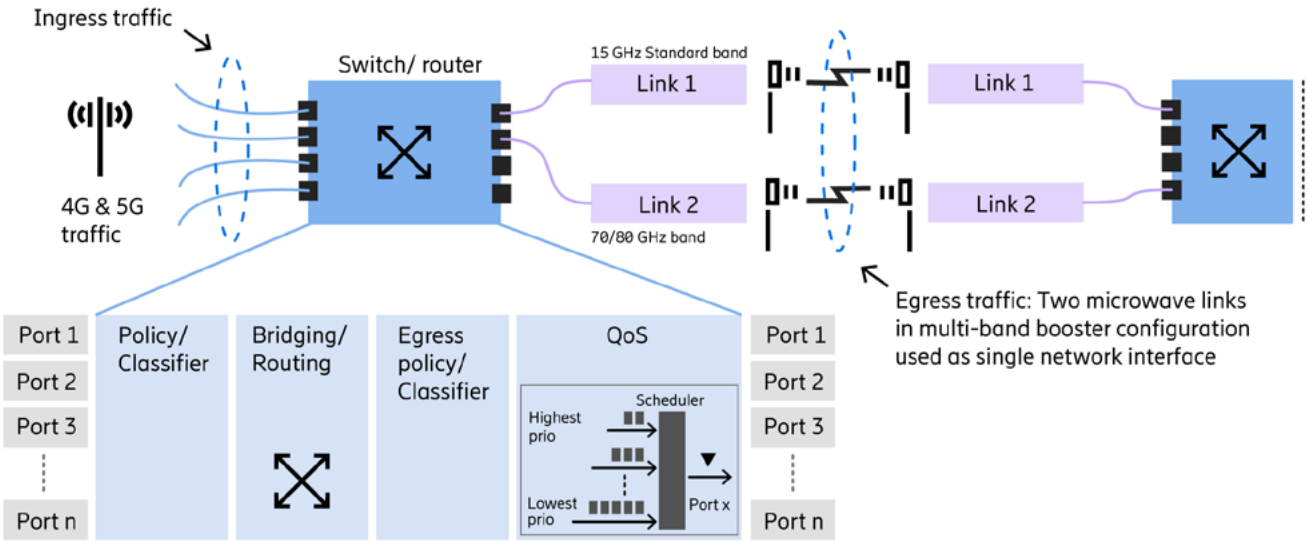


Figure 11. Aggregation test set-up

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.

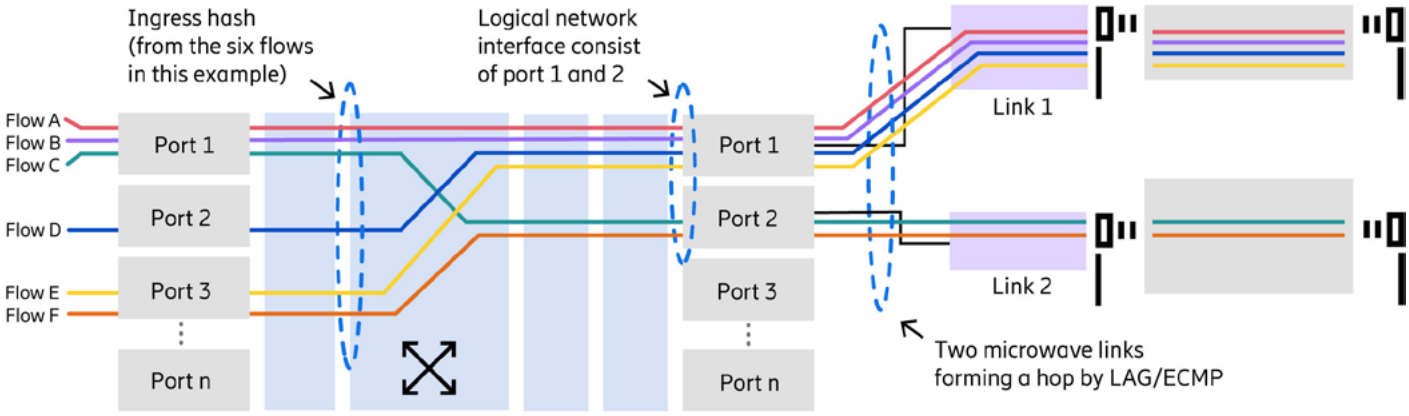


Figure 12. Link Aggregation (LAG) or Equal Cost Multi Path (ECMP) flows

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.

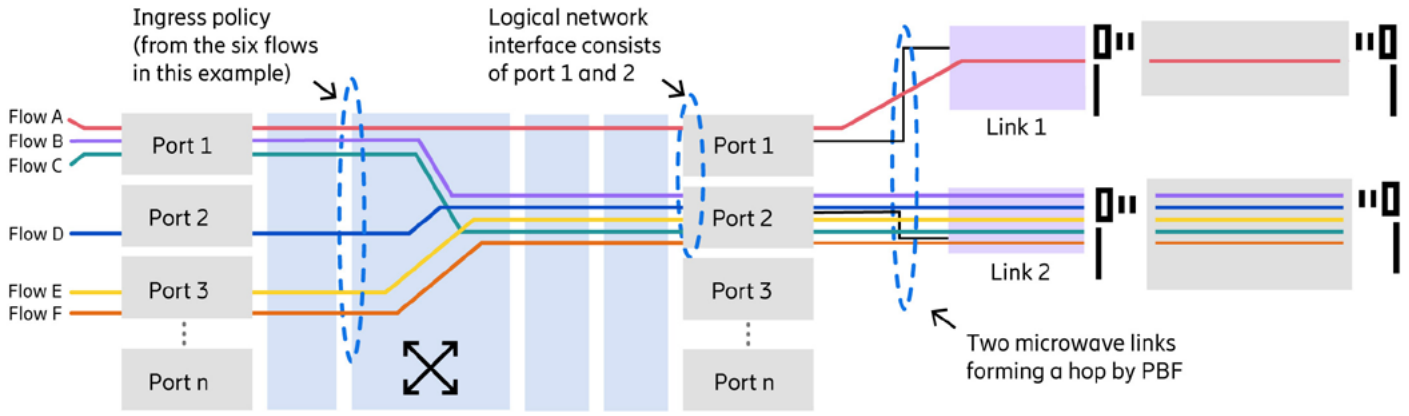


Figure 13. Policy based forwarding (PBF) flows

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.

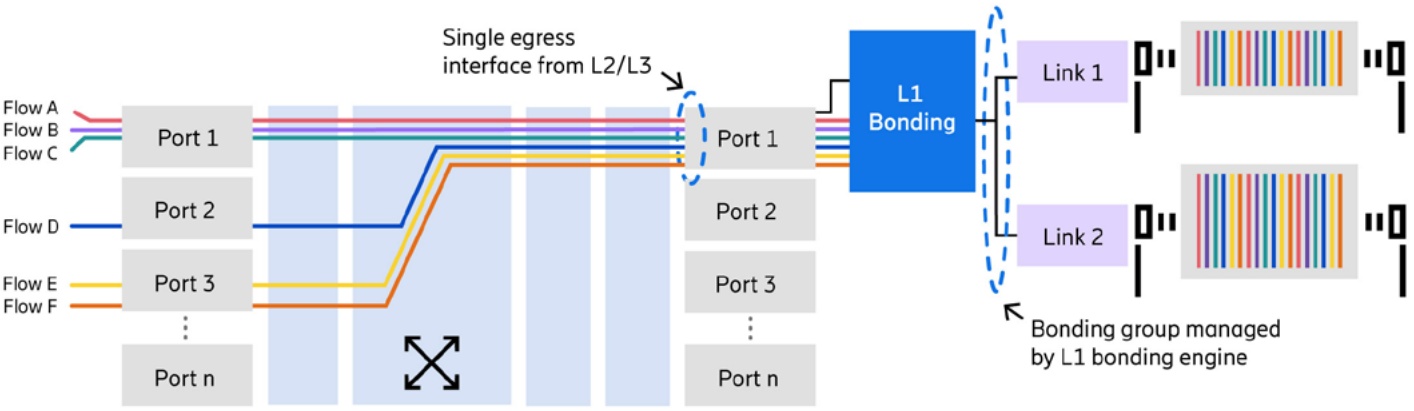


Figure 14. Layer 1 Bonding flows

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.

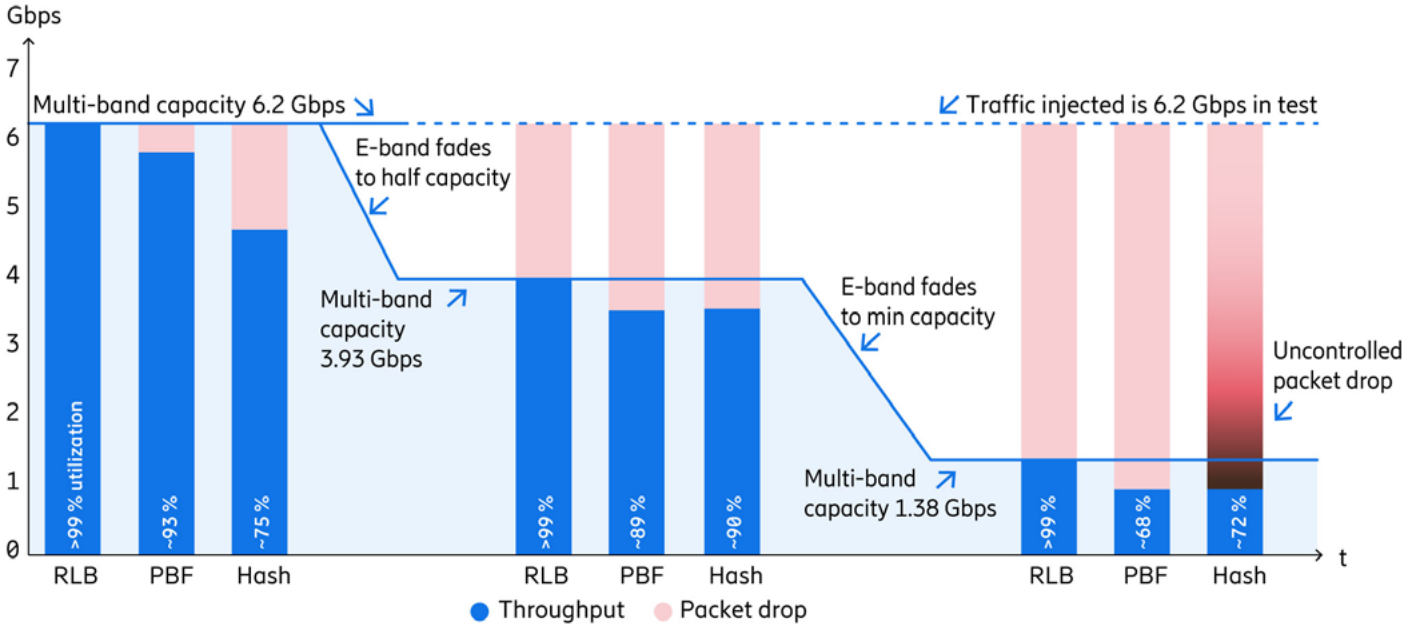


Figure 15. Throughput comparison during different fading events

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