

Hop Budget and Routing

How hop limits, routing protocols, and multi-hop link budgets determine end-to-end message delivery.

- [Understanding Hop Count and Hop Limits](#)
- [Designing for Multi-Hop Reliability](#)

Understanding Hop Count and Hop Limits

What Is a Hop?

A **hop** is a single radio transmission between two adjacent nodes. When a message originates at Node A and is received by Node B, that is one hop. If Node B re-transmits the message and Node C receives it, that is a second hop. Each re-transmission consumes airtime and contributes to channel congestion.

In a well-designed mesh, most messages should arrive in 1 - 3 hops. A message requiring 6 - 7 hops to reach its destination is a symptom of poor repeater placement (insufficient elevated infrastructure) or an oversized network diameter.

Hop Limits in Meshtastic

Meshtastic uses a **hop limit** field in the packet header to prevent infinite re-transmission loops. As of Meshtastic firmware 2.x (per meshtastic.org docs, verified 2026-06-08), the default is **3 hops** and the maximum configurable value is **7 hops**. Defaults can change between firmware versions, so verify against your firmware.

Each time a node re-transmits a packet, it decrements the hop limit by 1. When the hop limit reaches 0, the packet is not forwarded further. This is analogous to the TTL (Time To Live) field in IP networking.

```
Hop limit | Maximum relay chain
```

```
-----+-----
```

- 1 | Source → 1 relay → Destination
- 2 | Source → 2 relays → Destination
- 3 | Source → 3 relays → Destination (default)
- 5 | Source → 5 relays → Destination
- 7 | Source → 7 relays → Destination (maximum)

Note: the hop limit counts the number of re-transmissions (relays), not the total number of nodes in the chain. A packet with hop limit h undergoes up to h transmissions. The chain shown above is source + h relays + destination, where the destination is reached on the h -th (final) transmission and is itself the last node in the chain (so it does not re-transmit). A packet with hop limit 3 therefore reaches a destination up to 3 hops away.

Why Unlimited Hops Would Cause Broadcast Storms

Meshtastic uses flood routing: every relay node re-broadcasts every packet it receives. Without a hop limit, a packet could circulate indefinitely, with every node in the network repeatedly forwarding it. This creates a **broadcast storm** - a positive feedback loop where each re-broadcast triggers more re-broadcasts, rapidly consuming all available channel time and rendering the network unusable.

Packet deduplication (comparing packet IDs to avoid re-forwarding a packet already seen) and managed-flooding suppression (a node listens briefly before rebroadcasting and stays silent if it hears another node already relay the same packet), together with CSMA/CA contention windows, are the primary mechanisms preventing broadcast storms within a hop radius. The hop limit additionally bounds the maximum path length so a packet cannot circulate indefinitely.

Calculating Maximum Network Diameter from Hop Limit

With a hop limit of h , a packet undergoes up to h transmissions and traverses up to $h + 1$ nodes including the source, with the destination reached on the h -th hop. This means the maximum network "diameter" (the longest path between any two nodes) is h hops.

The figures below are **illustrative only**. Per-hop distance is highly terrain-dependent and is not a Meshtastic-published value: real LoRa links can be sub-1 km in dense urban areas or 100+ km in clear line-of-sight, depending on antenna height, spreading factor, transmit power, and obstructions. The table assumes a nominal ~ 5 km/hop (suburban) and ~ 15 km/hop (rural) purely to show how diameter scales with hop limit - do not treat these as guaranteed ranges.

Hop Limit	Illustrative path length (assumes ~5 km/hop, terrain-dependent)	Illustrative path length (assumes ~15 km/hop, rural, terrain-dependent)
1	5 km	15 km
3 (default)	15 km	45 km
5	25 km	75 km
7 (max)	35 km	105 km

For a well-designed urban network with elevated repeaters, a hop limit of 3 should be sufficient. If your network spans a wide rural area with widely spaced repeaters, you may need hop limit 4 or 5. Increasing beyond 5 is rarely justified and noticeably increases channel congestion.

Meshtastic Flood Routing vs. MeshCore Path-Based Routing

Meshtastic: Flood Routing

In Meshtastic, every node (by default, in CLIENT role) re-broadcasts packets it receives with a non-zero hop limit - flooding is not restricted to the router role. ROUTER/REPEATER roles have higher rebroadcast priority but are not required for flooding. The packet is sent once by each relay, regardless of whether a better path exists - though under managed flooding a relay suppresses its own rebroadcast if it first hears another node relay the same packet. This is simple and robust but inefficient: a 3-hop path in a dense network might result in dozens of redundant re-broadcasts from nodes along parallel paths.

Hop count behaviour: the hop limit is decremented at each relay. A packet leaving the source with hop limit 3 arrives at the destination (if 3 hops away) with hop limit 0. Packets consumed mid-network by deduplication still consumed their transmission slot on the channel.

MeshCore: Path-Based Routing

In MeshCore, the sender stores a learned path in its contact list. The path is embedded into the outgoing packet (source routing); each repeater forwards the packet only if its own address matches the next entry in the embedded path, otherwise it stays silent. Repeaters do *not* maintain a per-node routing table that they consult to pick a next hop. This is a form of source routing, where the sender specifies the path the packet should follow - not hop-by-hop layer-3 IP forwarding-table routing.

Hop count behaviour in MeshCore: MeshCore tracks hop and path information differently from Meshtastic's decrementing HopLimit TTL (the firmware has an internal maximum of 64 hops). This provides useful diagnostic information (you can see how many relays a received message used). MeshCore prevents routing loops through duplicate (packet) detection plus a hop limit: a packet cannot revisit a node (duplicate detection) and eventually expires regardless of network topology.

For network planners: in a Meshtastic network, channel load scales directly with hop limit and network density - every node hears and potentially re-transmits every packet. In MeshCore, channel load is more predictable because only the nodes on the actual path forward a given packet.

When to Increase Hop Limit - and the Tradeoffs

Consider increasing the hop limit above the default of 3 when:

- Your network covers a geographically large area (50+ km diameter) with widely spaced repeaters where 3 hops cannot span the network
- A specific critical path (e.g., a remote outpost to the network core) is consistently failing to deliver messages
- Path analysis shows that known routes require 4 - 5 hops due to terrain

Do not increase hop limit as a first response to poor coverage. Poor coverage is a placement problem, not a hop limit problem. Adding hops to compensate for missing repeaters increases channel congestion for all users without solving the root cause.

The tradeoffs of increasing hop limit:

Hop Limit	Network Reach	Channel Congestion	Typical Use Case
1	Very limited	Minimal	Dense urban, single repeater covers all clients
3 (default)	Moderate	Low - moderate	Well-planned city mesh with elevated anchor repeaters
5	Large	Moderate - high	Regional mesh spanning multiple cities or large rural area
7 (max)	Very large	High	Sparse wilderness mesh; use only when topology analysis confirms need

Designing for Multi-Hop Reliability

Link Budget Through Multiple Hops

In a multi-hop chain, each individual link (hop) must have a positive link margin. Unlike a wired network where signal is regenerated cleanly at each switch, a LoRa repeater decodes the incoming RF signal and re-transmits at full power, so the RF leaving each repeater is clean. But a packet that is only marginally decoded may contain bit errors or fail its CRC and be dropped — so each hop's received signal quality still determines whether the packet survives that hop. The re-transmitted signal is clean, but a weak receive at a repeater can still break the chain.

This is why **the weakest link in the chain determines end-to-end reliability**. If three hops each have reliabilities of 99%, 95%, and 90%, the end-to-end delivery probability is:

$$P(\text{delivery}) = 0.99 \times 0.95 \times 0.90 = 0.846 = 84.6\%$$

This is the key insight of multi-hop network design: each marginal link has an outsized negative effect on overall reliability because the probabilities multiply. A single 85% link with two near-perfect (~99%) links yields about $0.85 \times 0.99 \times 0.99 \approx 83\%$ end-to-end. If the other two links are also only moderate (~90% each), end-to-end delivery drops to roughly 70%.

Planning Hop Paths for Your Network

Before finalising repeater placement, map the expected relay chains for your most critical communication paths. The process:

1. **Draw the relay chain.** On your coverage map, trace the sequence of repeaters a message from a specific source node would traverse to reach the destination. In Meshtastic (flood routing), there may be multiple parallel paths; identify the primary path (shortest hop count with best margins) and the fallback path.

2. **Estimate link margin at each hop.** For each hop in the chain:
 - Calculate FSPL (Free Space Path Loss): $FSPL(dB) = 20 \cdot \log_{10}(d) + 20 \cdot \log_{10}(f) + 92.45$ (d in km, f in GHz)
 - Add terrain/vegetation loss (see Coverage Radius Estimation page for obstruction factors)
 - Add receiver antenna gain; subtract any receiver feedline/connector loss
 - Link Margin = TX Power + TX Antenna Gain – Path Loss + RX Antenna Gain – RX Feedline Loss – RX Sensitivity (all in dB, with TX power and sensitivity in dBm)
 3. **Identify where margin is thin.** A common planning rule of thumb is to flag any hop with less than 10 dB of margin as at-risk. (This is a heuristic, not a fixed standard — planners use anywhere from ~6 dB in benign environments to ~20 dB where high availability or heavy fading is expected.) Options: move the repeater to a better site, add a fill node on that hop, increase antenna gain on that link, or accept reduced reliability on that segment.
 4. **Document the analysis.** Record the estimated margin at each hop. Update after each wardriving survey that provides measured RSSI/SNR data on that link.
-

Worked Example: Three-Hop Path Analysis

Scenario: A remote ranch (Client A) communicates with the county EOC (Destination) via three repeaters (R1, R2, R3). Meshtastic on 915 MHz, SF10/BW125. TX power: 22 dBm (~160 mW) — the SX1262 chip maximum used by stock Meshtastic/MeshCore radios. (FCC Part 15.247 permits up to 30 dBm / 1 W conducted in 902 - 928 MHz, but reaching that requires an external power amplifier, and antenna gain above 6 dBi requires a 1-for-1 dB power reduction to stay within the 36 dBm EIRP limit. The figures below assume the buildable 22 dBm.) Antenna gain: 3 dBi omnidirectional on all nodes. Receiver sensitivity (SF10/BW125): –132 dBm.

Hop 1: Client A → Repeater R1 (hilltop, 4.2 km, rural open)

$$\begin{aligned}FSPL &= 20 \cdot \log_{10}(4.2) + 20 \cdot \log_{10}(0.915) + 92.45 \\ &= 12.46 + (-0.77) + 92.45 \\ &= 104.1 \text{ dB}\end{aligned}$$

Terrain loss (rural open): 0 dB additional (clear LOS)

$$\begin{aligned}\text{Link Margin} &= 22 \text{ (TX)} + 3 \text{ (TX ant)} - 104.1 \text{ (FSPL)} + 3 \text{ (RX ant)} - (-132) \text{ (sensitivity)} \\ &= 22 + 3 - 104.1 + 3 + 132 \\ &= 55.9 \text{ dB}\end{aligned}$$

Result: **55.9 dB margin** - excellent. This link is rock-solid.

Hop 2: Repeater R1 → Repeater R2 (rooftop, 9.8 km, suburban)

$$\begin{aligned}\text{FSPL} &= 20\text{-log}_{10}(9.8) + 20\text{-log}_{10}(0.915) + 92.45 \\ &= 19.82 + (-0.77) + 92.45 \\ &= 111.5 \text{ dB}\end{aligned}$$

Suburban obstruction add: +12 dB (unsourced rule-of-thumb clutter allowance; suburban clutter at 915 MHz commonly runs roughly +5 to +15 dB. For a rigorous estimate use a named model such as ITU-R P.1546 or COST-231 Hata.)

$$\begin{aligned}\text{Link Margin} &= 22 + 3 - (111.5 + 12) + 3 + 132 \\ &= 22 + 3 - 123.5 + 3 + 132 \\ &= 36.5 \text{ dB}\end{aligned}$$

Result: **36.5 dB margin** - good, well above the 10 dB minimum.

Hop 3: Repeater R2 → Destination EOC (18.5 km, suburban, EOC is in a building)

$$\begin{aligned}\text{FSPL} &= 20\text{-log}_{10}(18.5) + 20\text{-log}_{10}(0.915) + 92.45 \\ &= 25.34 + (-0.77) + 92.45 \\ &= 117.0 \text{ dB}\end{aligned}$$

Suburban obstruction add: +12 dB

Building penetration loss (EOC indoor): +10 dB (construction-dependent; building penetration at 915 MHz ranges roughly 5 to 25+ dB. See ITU-R P.2109.)

$$\begin{aligned}\text{Link Margin} &= 22 + 3 - (117.0 + 12 + 10) + 3 + 132 \\ &= 22 + 3 - 139.0 + 3 + 132 \\ &= 21.0 \text{ dB}\end{aligned}$$

Result: **21.0 dB margin** - acceptable (above 10 dB), but the EOC indoor penalty is significant. If the EOC uses a rooftop-mounted external antenna instead of an indoor unit, the 10 dB building penalty disappears and margin rises to about 31 dB. Strongly recommend external antenna at the EOC.

Chain Summary

Hop	Distance	Margin (dB)	Status
Client A → R1	4.2 km	55.9	Excellent
R1 → R2	9.8 km	36.5	Good
R2 → EOC	18.5 km	21.0	Marginal (address indoor loss)

End-to-end reliability of this chain is constrained by the R2 → EOC hop. Installing an external rooftop antenna at the EOC is the highest-priority action.

SNR and RSSI Thresholds for Reliable Forwarding

Theoretical link margins are estimates. In a live network, use measured RSSI and SNR to assess actual link quality. The thresholds below are approximate planning heuristics, not measured constants. They are also spreading-factor dependent: the figures here assume a mid-range preset (around SF9 - SF10). LoRa's SNR decode floor ranges from roughly -7.5 dB at SF7 down to about -20 dB at SF12, and at SF11/SF12 a link can decode reliably well below -120 dBm RSSI (sensitivity is around -131 dBm at SF11/250 kHz and -137 dBm at SF12/125 kHz). Treat the cutoffs as *margin relative to your preset's sensitivity* rather than fixed absolute numbers.

Metric	Minimum (marginal)	Target (reliable)	Good	Excellent
RSSI	-130 dBm	-120 dBm	-110 dBm	> -100 dBm
SNR	-15 dB	-10 dB	-5 dB	> 0 dB

Notes on these thresholds (approximate, and SF-dependent — assume roughly SF9 - SF10 here):

- **RSSI -130 dBm / SNR -15 dB:** A rough rule of thumb is a 50 - 70% success rate at a mid-range SF — packets may be decoded but unreliably. (At SF12 these levels are comfortably decodable.) Use only as an emergency fallback. Do not plan routes through links at this level.

- **RSSI –120 dBm / SNR –10 dB:** Practical minimum for planned routes at a mid-range SF. As a heuristic, expect roughly 85 - 95% packet delivery under normal conditions. Link will degrade in rain, vegetation growth, or when nearby interference increases the noise floor.
- **RSSI –110 dBm / SNR –5 dB:** Reliable for infrastructure links. Acceptable for primary repeater-to-repeater connections. Will maintain >98% delivery in most conditions.
- **RSSI > –100 dBm / SNR > 0 dB:** Strong link. Typical of well-placed nearby repeaters. These links rarely fail under normal operating conditions.

When reviewing live network telemetry, links consistently near the marginal end of the scale for your preset are candidates for remediation. Check the repeater placement, antenna alignment, and cable connections. If the physical setup is already optimal, a fill node on that path may be necessary.

The Weakest-Link Rule in Practice

When troubleshooting poor end-to-end delivery on a multi-hop path:

1. Collect per-hop SNR readings using the Meshtastic traceroute command (from firmware ≥ 2.5 it records the route back along with the Signal-to-Noise Ratio for each link — it reports SNR, not RSSI), or MeshCore path diagnostics. Per-link RSSI must be read locally from the node, not via traceroute.
2. Identify the hop with the lowest SNR (or locally-measured RSSI) - this is your weakest link.
3. Improving the weakest link will improve end-to-end delivery more than any other intervention. Do not chase marginal improvements on already-good hops.
4. After fixing the weakest link, re-test the full chain. A new weakest link may emerge.
5. Repeat until all hops meet your margin and SNR targets for the preset in use.