

MeshCore Repeaters

How to plan, configure, and deploy MeshCore repeaters to extend network coverage.

- [☰ Start Here — MeshCore Repeaters Guide](#)
- [Overview](#)
 - [What is a MeshCore Repeater?](#)
 - [Why Deploy a Repeater?](#)
- [Site Planning](#)
 - [Choosing a Repeater Location](#)
 - [Antenna Selection and Mounting](#)
- [Hardware](#)
 - [Hardware Considerations](#)
 - [Power and Solar Systems](#)
- [Configuration](#)
 - [Flashing Repeater Firmware](#)
 - [Radio Settings and Presets](#)
 - [Advertisements and Discovery](#)
 - [Terminal Chat CLI Commands](#)
- [Troubleshooting](#)
 - [Diagnosing MeshCore Repeater Problems](#)
 - [Reflashing and Factory Reset](#)
 - [Advanced MeshCore Repeater Diagnostics](#)
- [Advanced Repeater Configuration](#)
 - [MeshCore Repeater Name and Identity](#)
 - [Multi-Repeater Network Coordination](#)

- [MeshCore Repeater Diagnostics via Serial Console](#)
- [MeshCore Network Design](#)
 - [Planning a MeshCore Community Network](#)
 - [MeshCore vs Meshtastic: Choosing for Your Community](#)
 - [Repeater Density and Coverage Calculations](#)
- [MeshCore Repeater Hardware Builds](#)
 - [Budget MeshCore Repeater: Under \\$60](#)
 - [Pro MeshCore Solar Repeater: Complete Build](#)
- [MeshCore Network Expansion Strategies](#)
 - [When to Add a Repeater vs. When to Move One](#)
 - [Linking Isolated Mesh Islands](#)

☐☐ Start Here — MeshCore Repeaters Guide

This book covers everything about deploying MeshCore REPEATER firmware nodes - from choosing a site to building the hardware, configuring the firmware, and troubleshooting problems.

☐☐ Where to Start

1. [What is a MeshCore Repeater?](#) - Role in the network explained
2. [Choosing a Repeater Location](#) - Site selection principles
3. Then choose a build path: [Budget build \(~\\$60\)](#) or [Pro solar build \(~\\$217\)](#) (*approximate component totals as of June 2026; hardware prices vary by vendor and over time — check the linked build pages and current vendor listings before ordering*)

☐☐ What's In This Book

Planning and Site Selection

- [Choosing a Repeater Location](#)
- [Antenna Selection and Mounting](#)
- [Planning a MeshCore Community Network](#)
- [Repeater Density and Coverage Calculations](#)

Hardware Builds

- [Budget MeshCore Repeater: Under \\$60](#)
- [Pro MeshCore Solar Repeater: Complete Build](#)
- [Hardware Considerations](#)
- [Power and Solar Systems](#)

Configuration

- [Flashing Repeater Firmware](#)
- [Radio Settings and Presets](#)
- [Advertisements and Discovery](#)
- [MeshCore Repeater Name and Identity](#)
- [Multi-Repeater Network Coordination](#)

Network Design and Expansion

- [Room Servers & Gateways](#) - Placement strategy and adding a room server to your repeater site (covered in the dedicated Room Servers & Gateways book)
- [When to Add a Repeater vs. When to Move One](#)
- [Linking Isolated Mesh Islands](#)
- [MeshCore vs Meshtastic: Choosing for Your Community](#)

Troubleshooting

- [Diagnosing MeshCore Repeater Problems](#)
- [Advanced MeshCore Repeater Diagnostics](#)
- [Reflashing and Factory Reset](#)

➔ Related Books

- [MeshCore](#) - Protocol, firmware, and app guide
- [Room Servers & Gateways](#) - Adding a room server to your repeater site
- [Solar & Power Systems](#) - Powering outdoor repeaters
- [Antennas & RF](#) - Antenna selection and installation

Overview

What repeaters are, why they matter, and how they fit into a MeshCore network.

What is a MeshCore Repeater?

A MeshCore repeater is a device configured to run headlessly - without a connected phone or computer - whose sole job is to receive messages and forward them automatically. Repeaters form the backbone of any robust MeshCore network.

How a repeater works

Every MeshCore node listens for incoming radio packets. A repeater runs headless (no phone needed) and stores no message history, though some hardware (such as the Heltec V3 or T-Beam) does include a small status display. When it receives an eligible packet, it forwards it - skipping duplicates it has already seen and applying the network's forwarding rules - extending the message's reach to nodes that would otherwise be out of range.

This forwarding is automatic and requires no human intervention after initial deployment.

Why repeaters matter

Direct device-to-device range at ground level in a built-up area is often just a few hundred meters to roughly 1 km, depending heavily on terrain, obstructions, and antennas (range-test data for 915 MHz LoRa varies widely - as low as ~0.4 km in dense forest). A single repeater placed at elevation - a rooftop, hilltop, or communications tower - can extend coverage well beyond ground-level range: in favorable terrain with clear line of sight, often on the order of 10-20 km (6-12 miles) to nodes in its area, and farther node-to-node when both ends are elevated. Expect much less (a few km, or roughly 1-10 miles) where terrain or obstructions block line of sight. These are approximate, environment-dependent figures, not a guaranteed radius.

The more repeaters a community deploys in well-chosen locations, the more resilient and far-reaching the network becomes. Each additional repeater increases redundancy and reduces the chance of any single point of failure.

Repeater vs personal node

Personal node (BLE Companion)	Repeater
Paired with a smartphone	Runs standalone, no phone needed
Powered on/off by user	On continuously (solar or mains)
Used for sending/receiving messages	Forwards messages only
Typically carried or kept indoors	Mounted at elevation outdoors

Advertisements

Repeaters periodically broadcast **advertisements** announcing their presence on the network. These contain the repeater's identity, geographic position (if configured), and public encryption credentials. Other nodes use this information to discover the repeater and route messages through it.

As of current MeshCore firmware (2025-2026), the default flood advertisement interval for a repeater is 12 hours (set via `set flood.advert.interval <hours>`, range 3-168). Defaults can change between releases, so confirm yours with `get flood.advert.interval`. A zero-hop advert (`advert.zerohop`) is locally visible only; a flood advert (`advert`) propagates through other repeaters across the network.

Why Deploy a Repeater?

The case for community repeater infrastructure

A LoRa mesh network is only as strong as its infrastructure. Personal nodes carried in pockets or sitting in homes have limited range and go offline when their owners do. A well-placed repeater is always on, always forwarding, and serves every person in its coverage area simultaneously.

Key benefits

Extended range

A repeater at elevation can relay messages much farther than two ground-level handhelds could reach each other directly. With clear line of sight in favorable terrain, an elevated node can reach roughly 10 - 20 km (about 6 - 12 miles); a link between two elevated stations in ideal open terrain can stretch toward 20 - 25 miles, but expect substantially less with terrain obstruction or a handheld client. Without repeaters, two people a mile apart in a city might not be able to reach each other directly. Through a rooftop repeater, they can.

Network resilience

The more relay paths exist between any two points, the harder the network is to disrupt. If a node on a path goes offline, MeshCore can re-route — but only after the existing path fails and a new path is discovered, which introduces delay and can lose messages in the interim. Genuine redundancy requires a real alternate physical path, so build at least two independent routes for any critical link.

Always-on coverage

Unlike personal nodes that go offline when their owner's phone battery dies, a solar repeater operates indefinitely. Coverage is consistent regardless of whether individual users are active.

Multi-hop reach

MeshCore's firmware allows a high flood hop maximum (up to 64), but reliability falls off well before that. In practice, 3 - 5 hops through well-placed repeaters is the usable range and is enough to span considerable distances. A chain of rooftop or hilltop repeaters can cover an entire metro area or rural county.

Who should deploy a repeater?

Anyone with access to a good high location - a rooftop, a tall tree, a balcony with a clear view - can meaningfully contribute to local coverage. You do not need professional antenna installation experience for a basic install. However, outdoor and elevated mounts require proper weatherproofing, grounding and lightning protection (bond/ground per NEC 810 for fixed outdoor antennas), and care with any work at height. A simple pole mount and a weatherproof enclosure are often sufficient.

Site Planning

Choosing locations and understanding what affects coverage.

Choosing a Repeater Location

Location is the single most important factor in a repeater's effectiveness. A mediocre antenna on a perfect hilltop will outperform an excellent antenna at ground level every time.

Elevation is everything

At 915 MHz, radio waves travel essentially in straight lines (line-of-sight), with only limited diffraction around obstacles. The higher your repeater, the farther it can see before the curvature of the earth and obstacles block the signal. Greater height extends the radio horizon, though the gain per additional meter diminishes — range scales with the square root of height, so doubling antenna height does not double range.

Location types - best to acceptable

Hilltops and ridgelines

The gold standard. A repeater on a hilltop with unobstructed 360-degree views dramatically extends coverage beyond ground level — typically on the order of 10-20 km (6-12 miles) to handheld nodes, and farther node-to-node when both ends are elevated with clear line of sight. Longer links of 20+ miles are achievable only under ideal, full line-of-sight conditions between two well-elevated stations; that is a best-case radio horizon, not a guaranteed coverage radius to ground-level clients, and real coverage is usually less. Use a link-budget or viewshed tool to estimate realistically and verify with field tests. Even a modest hill of 100 - 200 feet above the surrounding terrain makes a significant difference.

Rooftops

Excellent for urban and suburban coverage. The highest accessible rooftop in a neighborhood, mounted on a pole or parapet wall, provides clear line-of-sight in most directions.

Communications and water towers

Already optimized for radio coverage. Many amateur radio operators and property owners are open to hosting community infrastructure. **Never climb a communications or water tower without the owner's explicit authorization** — doing so is often illegal, and such towers carry high-power RF emitters and electrical hazards. Arrange access and any mounting work through the tower owner and a qualified climber.

Tall trees

A practical option for rural or forested areas. **Safety first:** work at height on a tree or mast carries a serious fall risk — use proper fall-protection gear or hire a qualified climber, and never free-climb. Keep all masts and the antenna well clear of power lines. Mount as high as you can safely and legally reach, and ensure the solar panel receives adequate sunlight throughout the day.

Balconies and upper-floor windows

The minimum viable option when rooftop access is unavailable. Orient toward the direction offering the clearest line of sight. Even a second-floor position is meaningfully better than ground level.

What to avoid

- **Low ground** - valleys and depressions block signal in nearly all directions
- **Dense tree cover at antenna level** - foliage attenuates 915 MHz signals through absorption and scattering (roughly 0.2-0.5 dB/m of vegetation at 900 MHz, and more when wet); keep the antenna above the canopy
- **Large metal structures nearby** - HVAC equipment and metal roofing reflect and detune signals
- **Fully indoor placement** - walls absorb a significant fraction of signal strength

Testing before committing

Before a permanent installation, test with a temporary mount. Walk around the intended coverage area while watching signal on a paired phone. Tools like [HeyWhatsThat](#) can help visualize the theoretical radio horizon from a given point, though they do not account for buildings or vegetation.

Antenna Selection and Mounting

The antenna matters more than the radio

For a fixed repeater, the antenna is often the most impactful upgrade available. Moving from a 2 dBi stock antenna to a 6 dBi vertical on a rooftop pole adds about 4 dB of antenna gain, though some of that is offset by the feedline loss of the longer coax run a rooftop mount requires - so net gain at the antenna is usually less than 4 dB. The larger benefit is typically the improved line-of-sight from height, not the antenna gain alone. Note that higher gain comes with a narrower vertical pattern, which can create a dead zone directly below the antenna (see the omnidirectional notes below). **Mounting safety:** work on a roof or mast with fall protection, and keep the mast clear of overhead power lines by at least its full length plus a safety margin.

Antenna types for repeaters

Omnidirectional vertical (most common)

Radiates equally in all directions horizontally - ideal for a repeater that needs to serve a wide area. Higher gain (dBi) concentrates the signal closer to the horizontal plane, extending horizontal range but reducing coverage of areas directly below. As a general rule of thumb, moderate gain (roughly 3 - 6 dBi) is commonly a good balance for elevated repeaters; very high gain can starve coverage below the site. Very high gain antennas can create a dead zone directly beneath them.

Directional (Yagi, patch)

Focuses energy in one direction for maximum reach between two specific points. Requires careful aiming and is not suitable for general area coverage.

Gain vs. coverage angle

The figures below are general guidance, not fixed specifications - the right choice depends on your site, terrain, and mounting height.

Gain	Horizontal range	Best use
2 - 3 dBi	Short	Ground-level or indoor use
5 - 6 dBi	Medium - long	Most rooftop repeaters
8 - 9 dBi	Long, narrow beam	Tall towers over flat terrain (the narrow vertical beam needs height and level ground to avoid undershooting or overshooting on hills)

FCC compliance note: In the 902 - 928 MHz band, antennas above 6 dBi require an equal dB reduction in conducted transmit power under 47 CFR 15.247(b)(4). The maximum legal EIRP works out to 36 dBm (4 W). There is no 915 MHz point-to-point gain exemption (that applies only to 2.4 / 5.8 GHz). Plan power and antenna gain together.

Cable quality

Coaxial cable losses are significant at 915 MHz. RG-58 loses roughly 0.5 dB/m and even LMR-200 is about 0.33 dB/m - not truly low-loss for long runs. Keep the run as short as possible regardless of cable type. For runs beyond a few meters, use genuine low-loss cable in the LMR-240 / LMR-400 class (LMR-400 is about 0.11 dB/m); LMR-200 is acceptable only for very short runs. Weatherproof all outdoor connector joins with self-amalgamating tape or appropriate connector covers.

Key mounting rules

- Mount the antenna as high as practical, clear of obstructions in all directions. Use fall protection when working at height, and keep the mast clear of overhead power lines by its full length plus a margin.
- Keep the cable run short - locate the radio enclosure close to the antenna rather than running a long cable
- Use stainless steel hardware outdoors to prevent rust. Galvanic corrosion specifically occurs between dissimilar metals, so matching or electrically isolating the metals at a joint also matters.
- For a fixed outdoor antenna, bond and ground the installation per NEC Article 810.

- **Don't transmit without an antenna connected** - it's good practice to keep one attached. (The SX1262 used in these radios is fairly mismatch-tolerant, so a brief accidental transmit into an open connector is unlikely to destroy it; the risk of permanent damage from a mismatch is more relevant to high-power PA/FEM builds.)

Hardware

Power systems, enclosures, and hardware considerations.

Hardware Considerations

A MeshCore repeater needs three things: a LoRa radio running repeater firmware, an antenna, and reliable power. How you combine these depends on your deployment location and budget.

The LoRa radio

Any MeshCore-compatible LoRa device can be flashed with repeater firmware. The radio is rarely the performance bottleneck - location and antenna matter far more. Key requirements:

- **915 MHz band** - required for US/Canada. Beyond interoperability, the band choice is a legal one: 902-928 MHz is the FCC-authorized license-free ISM band in the US/Canada (47 CFR 15.247). The EU 868 MHz band is *not* authorized for this use in the US, so 868 MHz hardware (common in European product listings) should not be transmitted on here regardless of network compatibility — and it would not interoperate with the US network anyway.
- **External antenna connector** - essential for connecting a quality external antenna. Devices with only a PCB trace antenna are not suitable for fixed outdoor deployment.
- **MeshCore firmware compatibility** - verify against the MeshCore compatibility list before purchasing.

Purpose-built outdoor units vs. DIY

Purpose-built solar repeater units

Several manufacturers produce all-in-one weatherproof units with integrated solar panels, batteries, and LoRa radios. These are the simplest path to a permanent outdoor installation - they arrive ready to mount and flash.

Advantages: weatherproof from the factory, integrated power system, no enclosure engineering required.

Disadvantages: higher cost, limited hardware customization.

DIY builds

A builder can assemble a repeater from individual components: a LoRa board, weatherproof enclosure, solar panel, charge controller, and battery. The main challenges are reliable weatherproofing and correctly sized cable penetrations.

Advantages: full customization, potentially lower cost, complete control over every component.

Disadvantages: requires time and skill; waterproofing failure is a leading cause of field failures.

Enclosures

Electronics exposed to outdoor conditions should live in a weatherproof enclosure rated IP65 or higher. Note that the IP rating only holds if *every* penetration is sealed with a rated cable gland — drilling unsealed holes voids the rating. Key considerations:

- Proper cable glands on all penetrations (antenna, power, USB)
- Desiccant packs inside to absorb residual moisture
- UV-resistant material for sun exposure
- Thermal management - a sealed enclosure in direct sun can reach internal temperatures that exceed electronics and battery ratings (typically around 60 °C) without ventilation or shading. Shade the box or use a light-colored, UV-resistant material to reduce solar heating.

Power and Solar Systems

A repeater that runs out of power disappears from the network. Power system design is as critical as radio configuration for a reliable long-term deployment.

Why solar works for repeaters

MeshCore repeater firmware is designed for low power consumption. A repeater draws very little power when idle and slightly more when forwarding packets. This makes solar deployment practical even with modest hardware.

Sizing your power system

The goal: enough battery to run through several consecutive cloudy days, and a panel large enough to fully recharge on a typical sunny day.

- **Solar panel:** A 5 - 20W panel is reasonable example sizing for a low-draw repeater, but the right wattage depends on your load and your site's worst-month sun-hours. Mount it south-facing (in North America) and angle it roughly to match your latitude for best year-round output; if cloudy-season uptime is critical, tilt toward latitude +10-15 degrees to favor winter sun.
- **Battery chemistry:** LiFePO₄ (lithium iron phosphate) is strongly recommended for outdoor use. It tolerates cold *discharge* well, has a much longer cycle life than LiPo, and is significantly safer. **However, LiFePO₄ (like any lithium chemistry) must NOT be charged below 0 °C / 32 °F** - charging when frozen causes permanent lithium plating, reduced cycle life, and a fire risk. For cold climates, use a pack with a low-temperature charge cutoff in its BMS, or add a low-temp charge disconnect. Size for several days of runtime without any solar input - 3 - 5 days is a common minimum starting point; increase it for cloudy climates or critical links. For emergency-grade deployments, size the battery for your worst-case multi-day low-solar period (often longer than 3 - 5 days in winter) and validate it with a no-charge runtime test before relying on the node.
- **Charge controller:** Required between panel and battery. MPPT (Maximum Power Point Tracking) controllers are more efficient than PWM, especially in cold/temperate climates and on larger arrays. On very small systems the efficiency gain is modest, and a simple PWM controller is often sufficient and cheaper.

- **Fuse the battery.** Install an inline fuse in the battery positive lead close to the battery, sized per your controller and wiring. A LiFePO4 cell can deliver very high fault current; an unfused short can start a fire. Wire in order: battery → charge controller → panel.

Mains power

For rooftop installations with building power access, mains power plus a battery backup is more reliable than solar alone. Use a quality regulated supply and consider a small UPS to ride out brief power interruptions.

Power optimization

- Disable unused features: display backlight, Bluetooth, WiFi (if present on the board)
- Set a long **flood** advertisement interval for fixed infrastructure - the MeshCore default is commonly 12 hours, set via `set flood.advert.interval {hours}` (range ~3-168; verify in your firmware version). Note this is separate from the zero-hop `advert.interval` (default 0/off). More frequent ads increase power draw with minimal benefit.
- Do not set TX power higher than needed for your coverage goals - the power amplifier is the largest current draw during transmission. On a low-traffic repeater that spends most of its time receiving, idle/RX current may dominate total energy use, so also minimize wake/advertise frequency. TX power is also legally capped: 47 CFR 15.247 limits conducted power to 1 W (30 dBm) in 902-928 MHz, reduced dB-for-dB for antennas above 6 dBi.

Configuration

Flashing firmware, presets, radio settings, and advertisements.

Flashing Repeater Firmware

MeshCore-capable LoRa boards typically ship without MeshCore firmware (blank, vendor test firmware, or sometimes Meshtastic), so flashing MeshCore repeater firmware is required. This can be done entirely in a web browser - no software installation needed.

Using the MeshCore Web Flasher

1. Connect your device to your computer via USB
2. Open the [MeshCore Web Flasher](https://flasher.meshcore.io) (flasher.meshcore.io, as of June 2026) in Chrome or Edge (a Web Serial compatible browser is required; Firefox and Safari do not support it)
3. Select your device type from the list
4. Select the **Repeater** firmware variant
5. Click Flash and wait for the process to complete
6. The device will reboot automatically when done

After flashing

Once flashed, connect to the device using the MeshCore app (via Bluetooth) or the MeshCore serial console (via USB) to complete configuration. A newly flashed repeater is not inert - it boots on the firmware default frequency (869.525 MHz, the EU default) and will transmit on that default until changed. You **MUST** set the correct frequency/preset for your region before deploying, both for legal operation and so the repeater is visible to your network. For the USA and Canada, set the regional preset (910.525 MHz / SF7 / BW 62.5 kHz / CR 5).

Keeping firmware updated

MeshCore releases updates periodically with performance improvements and bug fixes. For a permanently deployed repeater, periodic firmware updates ensure you benefit from these improvements. The web flasher offers current official releases; select the latest stable build unless you have a reason to choose otherwise.

Radio Settings and Presets

Correct radio settings are essential for your repeater to interoperate with other nodes. The simplest and most reliable approach is to use the built-in preset.

Use the USA/Canada preset

In the MeshCore app, navigate to **Choose Preset** and select **USA/Canada (Recommended)**. This preset automatically applies the correct frequency plan, bandwidth, spreading factor, and coding rate for the North American MeshCore network.

“ **Do not manually set individual radio parameters unless you understand their effects and have a specific reason to deviate.** Incorrect settings will make your repeater invisible to the rest of the network.

What the preset sets

For reference, the USA/Canada preset resolves to settings within the 902 - 928 MHz band (commonly referred to as "915 MHz hardware"). As of the October 2025 "narrow" migration the US/Canada default preset is **910.525 MHz / SF7 / BW 62.5 kHz / CR 5**. These values can change with firmware updates, so the authoritative step is to read the actual values from the app or via the serial CLI (`get radio`) and match the network you intend to join.

Selecting the correct US/Canada region preset is also a compliance control: it confines the radio to the 902 - 928 MHz ISM band authorized under 47 CFR 15.247. Running a non-US preset (for example the EU 868 MHz default) on US hardware can place transmissions outside the authorized band.

Individual parameter reference

Parameter	Recommended	Notes
-----------	-------------	-------

Frequency	Set by preset	902 - 928 MHz (US ISM band). Do not use EU 868 MHz hardware on the US network.
Bandwidth (BW)	Preset default	Narrower BW = longer range, slower speed. Leave at preset default for network compatibility.
Spreading Factor (SF)	Preset default	Higher SF = better range, lower throughput. Keep at preset default.
Coding Rate (CR)	Preset default	Higher CR = better error correction, more overhead. Keep at preset default.
Flood hop maximum	Default (64)	MeshCore uses path-based routing, not a per-message hop-limit. The firmware flood hop maximum (<code>flood.max</code>) defaults to 64; you can cap flood propagation lower with <code>set flood.max <0-64></code> . In practice 3 - 5 hops spans substantial distances and delivery reliability falls off beyond that. The default is sufficient for most deployments.

Public vs. private channel

Configure your repeater on the **Public channel** so all network participants can read messages relayed through it. A repeater relays mesh traffic (flood adverts, public-channel messages, and routed direct messages) regardless of channel — it forwards packets without decrypting them. Private or custom channel keys only determine which nodes can decrypt those messages, not which repeaters forward them.

Advertisements and Discovery

MeshCore repeaters periodically broadcast **advertisements** - packets that announce the repeater's existence on the network. Other nodes use these advertisements to discover the repeater and include it in their routing decisions.

What an advertisement contains

- The repeater's identity (name and node ID)
- Geographic position (if configured - optional but useful for network mapping)
- Public encryption credentials used for routing - specifically, a public key identifying the repeater, which other nodes use to verify and route through it. This key is generated automatically; you do not configure it, and it is safe to broadcast.

Advertisement interval

The advertisement interval is configurable. The firmware supports a range of roughly 3 to 168 hours, with a default of 12 hours; set it with `set flood.advert.interval <hours>`. In observed community practice, deployments commonly use 6 - 12 hours for fixed infrastructure nodes. Longer intervals reduce radio traffic; shorter intervals help newly arrived nodes discover the repeater faster. Increasing frequency increases power consumption and adds traffic to the network with minimal benefit in a stable deployment.

Advertisement propagation: flood vs zero-hop

MeshCore advertisements use a binary propagation model - not a numeric hop count. You send either a flood advert or a zero-hop advert:

- **Zero-hop** (`advert.zerohop`): The advertisement is only visible to nodes within direct radio range. Useful for a local-only repeater.
- **Flood** (`advert`): The advertisement is propagated by other repeaters across the network, making the repeater discoverable to distant nodes. Use flood mode if you want your repeater to be visible across the wider network.

Setting a geographic position

Configuring your repeater's latitude and longitude allows it to appear on network maps and helps other operators understand coverage. Use your deployment location coordinates, not your home address. The position is included in flood advertisements and becomes visible to the wider network.

Terminal Chat CLI Commands

Terminal Chat CLI

Below are the commands you can enter into the Terminal Chat **client** (the companion/serial chat client). These are the Terminal Chat client commands, which differ from the repeater serial-console commands used elsewhere in this book. They are entered bare, with no host-tool prefix. Other pages that show `meshcli ...` are referring to a separate host-side tool with its own syntax - do not mix the two. For the on-device repeater serial CLI, see the canonical command reference at docs.meshcore.io/cli_commands.

```
set freq {frequency}
```

Set the LoRa frequency. Example: set freq 915.8

```
set tx {tx-power-dbm}
```

Sets LoRa transmit power in dBm.

```
set name {name}
```

Sets your advertisement name.

```
set lat {latitude}
```

Sets your advertisement map latitude. (decimal degrees)

```
set lon {longitude}
```

Sets your advertisement map longitude. (decimal degrees)

```
set dutycycle {percent}
```

Sets the transmit duty cycle limit (1-100%). Example: `set dutycycle 10` for 10%. Note: duty-cycle limiting is a regulatory requirement in the EU 868 MHz band (ETSI), not in the US 902-928 MHz band, where 47 CFR 15.247 uses frequency-hopping / digital-modulation rules instead. US operators are not subject to a duty-cycle cap, but courteous airtime limits still help network health.

```
set af {air-time-factor}
```

Sets the transmit air-time-factor. Deprecated - use `set dutycycle` instead.

```
time {epoch-secs}
```

Set the device clock using UNIX epoch seconds. Example: `time 1738242833`

```
advert
```

Sends an advertisement packet

```
clock
```

Displays current time per device's clock.

```
ver
```

Shows the device version and firmware build date.

```
card
```

Displays *your* 'business card', for other to manually `_import_`

```
import {card}
```

Imports the given card to your contacts.

```
list {n}
```

List all contacts by most recent. (optional {n}, is the last n by advertisement date)

```
to
```

Shows the name of current recipient contact. (for subsequent 'send' commands)

```
to {name-prefix}
```

Sets the recipient to the `_first_` matching contact (in 'list') by the name prefix. (ie. you don't have to type whole name)

```
send {text}
```

Sends the text message (as DM) to current recipient.

```
reset path
```

Resets the path to current recipient, for new path discovery.

```
public {text}
```

Sends the text message to the built-in 'public' group channel

Troubleshooting

Diagnosing MeshCore Repeater Problems

A systematic approach to diagnosing MeshCore repeater issues. Work from the physical layer up: radio first, then configuration, then network behavior.

Quick pre-diagnosis checklist

- Is the device powered? (check LED, screen, or measure supply voltage)
- Is an antenna connected? (transmitting without an antenna is poor practice and stresses the radio)
- Is the antenna connected to the LoRa SMA port, not the BLE u.FL? (Heltec V3/V4 have separate connectors)
- Is the device running repeater firmware, not Meshtastic or blank?
- Is it configured with the correct preset for your regional network?

Problem: Repeater is not appearing in MeshCore app contacts

Scenario A: Repeater has zero-hop advertisements

A repeater sending only zero-hop adverts (not flood adverts) is visible only to nodes within direct radio range. If you're testing from a distant location, it won't appear. Connect via USB serial and check the role and flood-advert configuration:

```
get role          # Confirm the node's configured role
get flood.advert.interval  # Flood advert cadence in hours (0 = zero-hop only)
```

To make the repeater visible network-wide, enable periodic flood adverts by setting a non-zero interval (default 12 hours), or send one immediately:

```
set flood.advert.interval 12  # Send a flood advert every 12 hours
advert                       # Send a flood advert immediately
```

Scenario B: Wrong preset

A repeater on a different preset than your phone is on cannot be heard. The phone and repeater must use identical preset settings (same SF, BW, CR, frequency). Verify both devices use the USA/Canada preset.

Scenario C: Not yet advertised

A freshly configured repeater won't be visible until it broadcasts its first advertisement. Default flood-advert interval is 12 hours. Force an immediate advertisement via the serial CLI:

```
advert # Serial CLI command to immediately broadcast a flood advertisement
```

Scenario D: Out of radio range

Test with a device you control. Bring it within 100 feet of the repeater (direct physical proximity, no obstacles). If it appears in contacts at that range, the repeater is working and the issue is range/placement. If it still doesn't appear at 100 feet, it's a configuration or hardware issue.

Problem: Repeater visible but messages not routing through it

Check: Repeater role set correctly

```
get role # Returns the node's configured role
```

A device running Companion (client) firmware will not relay messages for other nodes. To forward packets the device must run Repeater firmware (or Room Server firmware). The forwarding role is determined by which firmware variant you flash, not by a runtime command.

Check: Path discovery is completing

MeshCore must complete a route discovery (path discovery/acknowledgment exchange) before messages can be delivered. If messages are failing immediately after the repeater appears in contacts, path discovery may not have completed. Wait 60 seconds after first seeing the repeater, then try sending.

Check: Flood hop maximum

MeshCore uses path-based routing rather than a Meshtastic-style per-message hop-limit. The flood hop maximum (`flood.max`) defaults to 64, so there is no low default to "raise" — three to five hops is a practical reliability target, not a configured limit. If you have a specific reason to cap flood propagation, you can adjust it:

```
get flood.max          # Read the current flood hop maximum (default 64)
set flood.max 64      # Range 0-64
```

Problem: Repeater was working but stopped

Check: Power system

The most common cause of intermittent repeater failure. For solar-powered nodes:

- Battery voltage low: for a single LiFePO4 cell, ~3.2 V is the normal mid-charge plateau — concern begins below ~3.0 V and the cell is empty near 2.5 V. For a 4-cell (12 V) LiFePO4 pack, nominal is ~12.8 V and it is nearly discharged around 12.0 V. Check the solar panel and charge controller.
- Charge controller LED indicators: verify charging when sun is on the panel
- Winter: reduced daylight hours may be insufficient for the panel size. Consider adding battery capacity for winter.
- Snow or debris covering the panel: physically inspect

Check: Firmware crash or lock

Firmware can occasionally hang, especially after a power interruption during a write operation. A full power cycle (disconnect and reconnect power) usually resolves this. If the device is cycling rapidly (power LED flashing repeatedly), it may be boot-looping - likely a firmware corruption requiring reflash.

Check: Physical damage

Water intrusion is a common cause of permanent field failure. Signs: corrosion on the PCB, moisture beads inside the enclosure, condensation on components. Once water has infiltrated, the board often needs replacement. Prevention is much easier than cure - inspect enclosure seals annually.

Problem: Range is less than expected

1. **Antenna connected?** Transmitting into an open or disconnected SMA port causes VSWR problems and poor radiation.
2. **Correct antenna for 915 MHz?** An 868 MHz (EU-band) antenna on a US repeater will have higher SWR and reduced efficiency. US operation must stay within the 902-928 MHz ISM band under 47 CFR 15.247. Test with a NanoVNA (see the Antennas & RF section).
3. **Antenna mount location:** An antenna at 10 feet has dramatically less coverage than one at 30 feet. Every additional meter of height helps.
4. **Coax cable length and quality:** RG-58 loses roughly 0.5 dB/m at 915 MHz, so a 15-foot run loses about 2-2.5 dB and a ~50-foot run approaches 3 dB. Use LMR-200/240 or better and keep runs short.
5. **Obstructions at antenna level:** Metal roof edges, HVAC equipment, or tree branches within a few feet of the antenna scatter and attenuate signal significantly.
6. **TX power setting:** Verify TX power hasn't been accidentally lowered: `get tx`

Reflashing and Factory Reset

When a MeshCore repeater becomes unreachable, misconfigured, or boot-loops, reflashing is the recovery option. This page covers reflash procedures for all common device families.

When to reflash vs. when to reconfigure

Symptom	Try first	If that fails
Node working but wrong settings	Reconfigure via meshcore-cli or the serial CLI	Factory reset + reconfigure
Boot-looping or immediate crash	Power cycle; hold reset during boot	Reflash firmware
BLE not discoverable	Power cycle; V3: check BLE antenna	Reflash
Firmware corruption after interrupted flash	-	Reflash (required)
Unknown configuration state	Factory reset via serial: <code>erase</code> (destructive — wipes the device filesystem)	Reflash if serial inaccessible

Before reflashing: document your current config

If the device is still accessible, record its configuration before wiping. There is no single config-dump command in the MeshCore CLI — read each value with its own `get` command (typed into the serial console, or via the `meshcore-cli` host tool) and save the output:

```
# Capture each setting and save to a file
get name
get role
get freq
get radio
get tx
get lat
get lon
get flood.advert.interval
# Copy the console output into node_config_backup.txt
```

Reflash via MeshCore web flasher (recommended)

1. Connect device to PC via USB
2. Open flasher.meshcore.io in Chrome or Edge (WebSerial required)
3. Click "Connect" and select your device's serial port
4. Select device type and firmware variant (Repeater)
5. Click Flash and wait (~2 minutes)
6. Device reboots automatically when complete

Entering bootloader/DFU mode (when auto-detect fails)

ESP32-based devices (Heltec V3, V4, T-Beam)

Hold the BOOT button, press and release RESET, then release BOOT. The device enters download mode and should appear as a serial port. Some devices require the BOOT button held while inserting the USB cable.

nRF52840-based devices (RAK4631, T-Echo, Nano G2 Ultra)

Double-tap the reset button quickly. The device enters DFU mode and appears as a USB drive named "BOOT" or similar. Copy the .uf2 firmware file to this drive to flash.

Alternatively, use the web flasher which handles DFU automatically for most nRF52840 devices.

Station G2

Power requirement: The Station G2 generally expects a higher input than a standard 5V USB cable (commonly a 9–19V DC or USB-C PD source). Confirm the exact input-voltage range against the manufacturer's datasheet for your unit before powering it — feeding the wrong voltage can destroy hardware. If a board fails to flash on a plain 5V cable, an under-powered supply is a likely cause.

Post-reflash configuration

After flashing, the device has factory defaults. Note that the repeater **role is determined by the firmware variant you flashed** (flash the Repeater build) — there is no command to change the role; you can only read it with `get role`. Reconfigure the radio and identity using the serial console (or `meshcore-cli`):

```
# Apply the USA/Canada radio plan (910.525 MHz / BW 62.5 kHz / SF7 / CR5).
# You can select the preset in the MeshCore app, or set the raw params directly:
set radio 910.525,62.5,7,5
reboot

# Confirm the role (set by the flashed firmware, read-only):
get role

# Set name (use your documented name)
set name "REPEATER-NAME"

# Set position (latitude and longitude are separate commands;
# there is no altitude setting in the firmware):
set lat 47.6062
```

```
set lon -122.3321

# Advertisements: send one flood advert now, and set the periodic
# flood advert interval in HOURS (default 12, range 3-168):
advert
set flood.advert.interval 12

# Set TX power in dBm. Real command is "set tx" (documented range 1-22 dBm;
# the SX1262 PA tops out near 22 dBm). 22 matches the build guides.
# Total EIRP = TX power + antenna gain - feedline loss; keep it within the
# FCC 47 CFR 15.247 limits for 902-928 MHz. With an antenna over 6 dBi you
# must reduce conducted power dB-for-dB, so lower this below 22 as gain rises.
set tx 22

# Verify the key settings:
get role
get radio
get freq
get tx
get name

# Reboot
reboot
```

When the device is completely unresponsive

If USB serial doesn't appear and the device shows no activity:

1. **Try a different USB cable.** Many micro-USB and USB-C cables are charge-only and have no data lines. Use a known-good data cable.
2. **Try a different USB port.** USB hubs can cause issues; try a direct port on the computer.
3. **Check the USB driver if needed:** The Heltec V4 (ESP32-S3) uses native USB and does *not* require a CH340/CP210x bridge driver — it enumerates as a standard USB CDC serial device on Windows 10/11. (The V3 used a CP2102 bridge.) RAK4631 and T-Echo also use standard USB CDC. If a unit fails to enumerate, check the cable and DFU/download mode rather than a bridge driver.

4. **Check for physical damage:** Inspect the USB port for bent pins or corrosion. A damaged USB port prevents flashing.
5. **Last resort:** Some boards can be recovered via JTAG/SWD with a debug probe. Consult the manufacturer's documentation.

Advanced MeshCore Repeater Diagnostics

When basic troubleshooting hasn't resolved a repeater issue, these advanced diagnostic techniques help identify hardware failures, RF problems, and protocol-level issues.

Systematic RF Path Verification

Before concluding a repeater has failed, verify the RF path independently:

1. **Local RF check:** Bring a test node within 1 meter of the suspect repeater. Can it hear the repeater's transmissions? If yes, the repeater is transmitting. If no, the repeater may not be transmitting.
2. **Antenna check:** Disconnect the antenna from the repeater and check for physical damage (bent connector, water ingress at connector). Re-connect securely.
3. **Cable check:** If using a remote antenna with coax, check the cable for damage, kinks, or water ingress at connectors. Substituting a known-good short cable bypasses cable issues.
4. **SWR check:** If you have a NanoVNA, check antenna SWR. A sudden SWR change vs. baseline indicates antenna or cable change (ice loading, lightning damage, physical damage).

Firmware Diagnostic Commands

These commands are entered over the device's serial/terminal CLI (the same console used elsewhere in this book). The authoritative command reference is docs.meshcore.io/cli_commands — the commands below match it. A MeshCore repeater runs firmware on a microcontroller (nRF52840 or ESP32); there is no Linux shell or filesystem on the device, so only these serial CLI commands work.

```
# Show packet counters (received / sent):  
stats-packets  
  
# Show radio statistics (RSSI, SNR, airtime, RX errors):
```

```
stats-radio

# List nearby neighbors heard via advert, with SNR.
# Output is the 8 most recent adverts, each line encoded as
# {pubkey-prefix}:{timestamp}:{snr*4} (SNR only, not RSSI):
neighbors

# Show core stats (battery, uptime, queue length, debug flags):
stats-core

# Print the captured RX log to the serial terminal.
# Use "log start" / "log stop" to control capture, then read it
# back and look for error lines:
log
```

MeshCore has no serial `ping` command and no per-node forced route-discovery command. To exercise routing to a specific destination, send a message to that contact from a client and observe whether a path establishes. Zero-hop neighbor discovery (on firmware that supports it) is `discover.neighbors`.

Power System Deep Diagnostics

Solar-powered repeaters frequently experience issues traced to power rather than RF:

```
# Check battery, uptime and queue length over the serial CLI:
stats-core

# On boards with power management, the last reset cause and boot
# voltage can be read (where supported by your firmware build):
get pwrmtg.bootreason
get pwrmtg.bootmv

# Low battery can cause a brownout during TX: the TX current spike
# drops voltage below the regulator minimum and the board resets.
# There is no Linux log file on the device to read; use stats-core
# (or an external voltage logger) to track voltage history.
```

Signs of power-related resets:

- Repeater uptime counter resets during evening/night hours (low SOC)
- Resets correlate with high-traffic periods (many TX events = high current draw)
- Stable operation after adding a larger battery or bypass capacitor at the node power input

Isolating Radio Hardware Failure

If you suspect the SX1262 radio chip has failed:

- **Verify SPI communication:** Most MeshCore firmware reports a radio-init status line on the serial boot log. A successful SX1262 init is logged; an init-failure line or no radio status at all suggests a hardware problem. The exact wording varies by firmware version — check your boot log rather than searching for one fixed string.
- **Substitution test:** Swap the suspect board with a known-good board. If the problem follows the board, it's a hardware failure.
- **Visual inspection:** Check for cold solder joints on the LoRa module, corrosion on SPI pins, or physical damage to the radio module.

When to Declare Hardware Dead

Replace the hardware rather than continuing to debug when:

- The board fails to initialize the radio after firmware reflash
- The board has confirmed water ingress damage (corrosion on PCB, discolored/burned SMD components)
- The board passes diagnostics but still fails in production consistently after replacing all external factors (antenna, cable, power)
- The cost of continued debug time exceeds the cost of a replacement board (\$30-60)

Advanced Repeater Configuration

MeshCore Repeater Name and Identity

Every MeshCore repeater broadcasts an identity advertisement that makes it visible to the network. Setting a meaningful name and position makes your repeater useful to the wider community.

Setting a Repeater Name

The repeater name appears in client apps when users are selecting which repeaters to contact or route through. Use a consistent naming convention for community repeaters:

- **Geographic convention** - Name by location: `W5ABC-Mt-Wilson`, `K6XYZ-Oakland-Hills`, `WA7QRS-Snoqualmie`
- **Grid square convention** - Used by some communities: `DM04-Repeater-1`
- **Descriptive** - Simple but clear: `Downtown-SF-Roof`, `I-5-Corridor-N`

Note: Call signs in these examples are just a human naming convention. MeshCore devices operate under FCC Part 15 (license-free) on 902-928 MHz, not Part 97 amateur rules - no station ID is required, and a call sign in the name does not make the device an amateur-radio station.

Configure via the MeshCore app or serial console:

```
set name MyRepeaterName
```

The node name maximum length is 24 bytes if a location is set, 32 bytes otherwise (the limit is measured in bytes, and emoji or accented characters use more than one byte each). Plain ASCII names of roughly 24 characters or fewer are safe. As a separate style suggestion, keeping names short also improves display readability across client apps.

Setting Geographic Position

Position data is included in flood advertisements, making your repeater visible on network maps and helping operators understand coverage. Configure with decimal degrees:

```
set lat 37.7749
set lon -122.4194
```

MeshCore position advertisements use latitude and longitude only; there is no altitude setting in the CLI. Use your deployment coordinates, not your home address - the position is broadcast to the entire network.

Advertisement Configuration

Control how often your repeater sends its flood advertisement:

```
set flood.advert.interval 12
```

The flood advert interval is specified in hours (range 3-168), with a repeater default of 12. The 12-hour default is appropriate for stable, permanent deployments. To verify a repeater is appearing in client apps during initial deployment, temporarily lower the flood interval (for example to 3 hours) and then return it to 12 for steady-state operation to minimize network load. A separate zero-hop advert interval, `set advert.interval <minutes>` (60-240 minutes), controls only the local zero-hop advert.

Flood vs Local Advertisement

MeshCore adverts are sent as one of two types - there is no per-advert hop-count setting. To send a flood (multi-hop) advert that other repeaters propagate across the network, run:

```
advert
```

`advert.zerohop` - sends a zero-hop advert visible only to nodes in direct radio range (local-only).

`advert` - sends a flood advert that propagates through the mesh, making the repeater discoverable to distant nodes; its propagation depth is limited by `flood.max` (range 0-64, default 64), not by a per-advert hop count. Use flood adverts (`advert`) for public community repeaters; use zero-hop adverts (`advert.zerohop`) for private or testing nodes.

Multi-Repeater Network Coordination

When multiple MeshCore repeaters serve the same community, coordination between operators ensures the network behaves predictably and provides maximum benefit to users.

Radio Preset Consistency (and the Channel-Key Myth)

A common misconception is that every repeater on a community network must share a "channel key." It does not. In MeshCore, **repeaters forward packets based on routing/path information and do not decrypt message payloads** — they relay encrypted traffic for any channel without needing that channel's key. Channel keys are a client-side concern: only clients holding a channel's key can read that channel's group messages. So a repeater never "silently drops traffic it cannot decrypt." What every repeater (and client) *must* share to interoperate is the same **radio preset / frequency** — if the radio parameters don't match, the repeater literally cannot hear the packets. When deploying a new repeater, coordinate the radio preset, frequency, and repeater naming/placement with the network coordinator before commissioning.

Verify a new repeater is relaying by sending a test message from a client node so that it routes through the new repeater. A successful delivery confirms the repeater is hearing and forwarding correctly (correct preset, placement, and RF range). It does *not* test channel-key alignment — that is verified between client nodes (whether two clients can read each other's group messages), not at the repeater. A failed route-through test points to a radio-config/preset mismatch or insufficient RF range, not a missing channel key.

Coverage Overlap Planning

Adjacent repeaters should have some coverage overlap for redundancy and continuity. The exact amount depends on terrain, antenna height, and node density — there is no universal percentage. Validate overlap empirically with drive/walk tests confirming each area is reachable via at least two repeaters. Overlap provides:

- **Redundancy** - if one repeater goes offline, adjacent repeaters still serve the overlap area
- **Continuity of coverage for moving nodes** - when a mobile node leaves one repeater's range, its cached path will fail and MeshCore re-runs path discovery to find a route through the next in-range repeater. This is not a seamless cellular-style handoff — there is no association handed between repeaters; any in-range repeater simply relays. Expect a brief delivery gap while a new path is discovered. Overlap reduces, but does not eliminate, this gap.
- **Route diversity** - multiple path options between distant network segments improve end-to-end reliability

Insufficient overlap creates coverage holes where users are out of range of all repeaters. At the other extreme, deploying many repeaters very close together — covering the same footprint with no new coverage — adds channel airtime contention and flood amplification without proportional benefit. Note that close spacing is not inherently wasteful: in dense urban terrain with heavy building attenuation, sub-kilometer spacing can be necessary and appropriate (the density guidance suggests roughly one repeater per $\sim 1 \text{ km}^2$ / $\sim 500 \text{ m}$ radius). Avoid overlap only where it is genuinely redundant.

Frequency and Preset Coordination

All community repeaters must use the same radio parameters (the USA/Canada preset is recommended in this region). Verify with:

```
get radio
```

The `get radio` output shows the current frequency, bandwidth, spreading factor, and coding rate (`get freq` and `get tx` report frequency and TX power individually). MeshCore presets are an app-side selection; the serial CLI reports the raw radio parameters rather than a named preset. Confirm these values match the community-standard USA/Canada preset before committing a new repeater to service.

Network Documentation

Maintain a community document with:

- Repeater name, operator callsign/contact
- Physical location (approximate - GPS coordinates if the operator consents to publishing)
- Power system type (mains, solar) and expected uptime
- Installation date and last maintenance date
- Known coverage limitations or issues

This documentation is invaluable when diagnosing network problems or planning expansion. Store it in a shared document that all operators can access and update.

Repeater Retirement and Replacement

When a repeater is permanently taken offline, notify the community so they can update routing expectations and coverage maps. Removing a node that other nodes' cached routes depend on will cause temporary routing failures until routes are rediscovered. This is normal behavior; MeshCore re-discovers routes when existing paths fail.

MeshCore Repeater

Diagnostics via Serial Console

The MeshCore serial console provides direct access to repeater state and diagnostic information. Connecting via USB to a deployed repeater is the most reliable way to diagnose problems that cannot be addressed remotely.

Connecting to the Serial Console

On Windows: use PuTTY or the Arduino Serial Monitor. On Linux/Mac: use `screen` or `minicom` as a serial terminal.

```
# Linux/Mac
screen /dev/ttyUSB0 115200

# Windows (PuTTY): Connection Type = Serial, Speed = 115200, COM port varies
```

MeshCore boards use 115200 baud for the serial console, including RAK boards. (If you have configured an attached GPS module, that module's own UART may run at a different baud rate such as 9600 - but that is a separate interface from the console.)

Key Diagnostic Commands

These are commands from the on-device MeshCore serial CLI (see docs.meshcore.io/cli_commands). There is no single combined "status" command - health information is split across several `get ...` and `stats-...` commands.

Command	Output	Use for
<code>ver</code>	Firmware version	Confirm firmware build

Command	Output	Use for
<code>stats-core</code>	Battery voltage, uptime, queue depth	Overall health check
<code>get radio</code> / <code>get freq</code>	Radio config (frequency, bandwidth, SF, coding rate)	Verify radio settings match the network
<code>neighbors</code>	Up to the 8 most recently heard nodes, with timestamp and SNR	Verify which nodes are reaching this repeater
<code>stats-packets</code>	Packet counters: Received and Sent	Identify traffic/routing problems
<code>stats-radio</code>	Noise floor, last RSSI/SNR, airtime, receive errors	Signal quality of last received packet
<code>log</code>	Prints a captured rx log (must first be started with <code>log start</code> , stopped with <code>log stop</code> , cleared with <code>log erase</code>)	Capture and review received-packet activity
<code>reboot</code>	(restarts device)	Recover from hung state

Note: `contacts` and `list` are companion/client-side concepts, not repeater serial commands - on a headless repeater, use `neighbors` instead. The `log` command is a packet/rx capture you start and stop, not an always-on event log.

Interpreting Stats Output

The `stats-packets` command is a useful diagnostic tool. The firmware exposes Received and Sent counters. A healthy repeater shows:

- **Sent count tracking received traffic** - the repeater is relaying packets it is meant to forward. (The documented counters are Received and Sent; there is no separate "forwarded" or "dropped" counter, so judge activity from the Sent counter rising alongside Received rather than from a "forward rate" metric.)
- **Drops in MeshCore** - MeshCore drops a flood packet when it exceeds the `flood.max` hop ceiling (default 64), and its loop detection (`loop.detect`) drops a packet whose path shows this repeater's own ID/hash repeated. This is not a Meshtastic-style per-packet hop counter decrementing to zero. In MeshCore's path-based routing, a repeater also intentionally does *not* retransmit packets whose embedded path does not include it - these appear as non-forwarded traffic but are correct behavior, not a loop. Distinguish this selective non-forwarding from a genuine duplicate-flood loop before acting.
- **Increasing received count over time** - Confirms the repeater is hearing traffic from the network

Common Issues and Diagnostics

Symptom	Check	Fix
Received count stays at zero	Check antenna connection, verify radio settings match network	Reconnect antenna; verify radio parameters with <code>get radio</code> and <code>get freq</code> (frequency, bandwidth, SF, coding rate) and confirm they match the USA/Canada preset values shown in the app
Sent count zero despite receives	Verify device is running repeater firmware variant and that forwarding is enabled	Reflash with repeater firmware; confirm <code>set repeat on</code>
Battery voltage declining	Check solar panel output, charge controller LVD setting	Clean panel, verify charge controller settings
Rebooting frequently	Check for low battery voltage causing brownout	Size battery correctly; check charge controller
Not appearing in client node list	Repeater may not be sending flood adverts; check the flood advert interval with <code>get flood.advert.interval</code>	Send a flood advertisement with <code>advert</code> and confirm a flood advert interval is set, e.g. <code>set flood.advert.interval 12</code> (hours). Zero-hop adverts (<code>advert.zerohop / set advert.interval</code>) are local only. There is no <code>advert_hops</code> parameter.

MeshCore Network Design

Planning a MeshCore Community Network

Deploying a MeshCore network for a community requires planning beyond simply placing repeaters - you need to think about coverage, redundancy, operator coordination, and long-term maintenance.

Phase 1: Define Coverage Goals

Before placing a single node, answer these questions:

- **Who are the users?** - Community members, emergency responders, ARES team, public? Each has different requirements.
- **What is the geographic target area?** - City, county, neighborhood, trail corridor? Define the boundary on a map.
- **What quality of service is needed?** - Best-effort casual use vs. mission-critical emergency communications have different reliability requirements.
- **What is the budget?** - A hobbyist community network might rely on volunteer-hosted nodes; a professional deployment might have a funded infrastructure budget.

Phase 2: Identify Candidate Sites

For each site, evaluate:

1. **Elevation** - Higher is almost always better. Use topographic maps or USGS terrain data to identify hilltops, water towers, or tall buildings that have commanding views of the target area.
2. **Power availability** - Mains power is most reliable; solar works at most outdoor sites; battery-only is acceptable for temporary or low-priority nodes.
3. **Accessibility** - You will need to access this site for installation and maintenance. A perfect hilltop that requires technical climbing is impractical for most operators - and it is a real safety and liability hazard, not just an inconvenience. High sites that require climbing carry genuine fall and legal risk: use professionals for tower work, get written

permission, and never climb without proper fall-protection gear and authorization.

4. **Permission** - Property owner permission is required. Start with sites where you have existing relationships: your own roof, a cooperating business, a friendly landowner.

Phase 3: Coverage Analysis

Before installing, verify coverage using tools:

- **HeyWhatsThat.com** - Enter your proposed repeater coordinates and height; generates a viewshed map showing where has line-of-sight to that point.
- **Radio Mobile Online** - More detailed propagation modeling including terrain, frequency, and antenna parameters.
- **Field testing** - Place a temporary node at the candidate site and drive/walk the intended coverage area while monitoring packet reception. This is ground truth.

Phase 4: Deployment Sequence

1. **Deploy the backbone first** - Install your 2-3 highest-coverage repeaters before deploying fill nodes. The backbone provides the largest coverage gain per node installed.
2. **Test between backbone nodes** - Verify each backbone node can reach at least TWO other backbone nodes by independent paths, so the loss of any single node does not partition the network. Avoid daisy-chain topologies for emergency-grade networks - a chain that looks tested can still split in two when one node drops.
3. **Add fill nodes for dead zones** - Use coverage testing to identify gaps; deploy targeted fill repeaters.
4. **Recruit community members** - Once basic coverage exists, recruit nearby property owners to host additional nodes. Their rooftops fill gaps and add redundancy.

Documentation and Handoff

Create documentation before deploying each node:

- Physical address or GPS coordinates of the site
- Property owner name and contact information
- Equipment installed (board type, firmware version, power system)
- Configuration (name, channel, advertisement settings)
- Access procedure (how to reach the installation for maintenance)
- Emergency contact (who to call if the node is causing problems)

Store this documentation somewhere accessible to all network operators - not just one person's laptop.

MeshCore vs Meshtastic: Choosing for Your Community

If you're building a community mesh from scratch, choosing between MeshCore and Meshtastic is one of the first decisions. This page provides a framework for that decision.

The Most Important Factor: Community

The single most important factor is what your local community already uses. A technically inferior protocol with 50 active nodes in your area is better than the technically superior protocol with zero. **Check what's already deployed in your area before committing.**

- For MeshCore nodes, check the MeshCore map at meshcore.co.uk/map.html. Note that meshmap.net maps Meshtastic nodes only - it will not show MeshCore presence, so don't rely on it to gauge MeshCore adoption (as of 2026; verify the current map URLs).
- Search for local ham radio ARES/EMCOMM groups - many have adopted one protocol
- Ask in local ham radio clubs and maker communities, and ask local MeshCore operator groups directly

Choosing MeshCore When

- Your area already has MeshCore infrastructure or a MeshCore operator community
- You are building dedicated repeater infrastructure for a larger network (50+ nodes)
- You want public-key encrypted DMs. Both projects now use Curve25519 ECDH for direct messages (Meshtastic since v2.5, released 2024); neither MeshCore nor Meshtastic provides forward secrecy. MeshCore is not uniquely strong here.

- You have technically sophisticated operators who understand routing and can configure path-based routing
- You are building a network where room servers and internet connectivity are part of the design

Choosing Meshtastic When

- Your area already has an active Meshtastic community
- You want the widest hardware compatibility and largest ecosystem
- Your user base is non-technical and needs the most polished, easy-to-use apps
- You want the most beginner-friendly experience for recruiting new members
- Your network is small enough that managed flooding works well. Meshtastic does not publish a hard node-count limit - scaling depends on traffic, airtime, and the configured hop limit, not a fixed node count.

Running Both

Some communities operate parallel Meshtastic and MeshCore networks. This is common in areas where early adopters chose different protocols. The networks operate on the same frequency band but use different packet formats and cannot interoperate. A single operator can run both by using two LoRa boards - one flashed as MeshCore, one as Meshtastic.

Running parallel networks adds complexity but ensures coverage for all community members regardless of which protocol they chose. If your community has both, coordinate channel settings and coverage to complement rather than duplicate each other.

Summary Comparison Table

Factor	MeshCore	Meshtastic
Routing	Path-based (path discovery/acknowledgment)	Flooding
Scales to	Larger networks (path-based routing scales better than flooding)	Smaller networks; flooding degrades as traffic grows. Exact node counts depend on traffic, airtime, and topology - load-test your own deployment.
DM encryption	Curve25519 ECDH + AES-128 (no forward secrecy)	PSK pre-v2.5; Curve25519 ECDH since v2.5 (no forward secrecy)

Factor	MeshCore	Meshtastic
App ecosystem	Smaller	Larger (Android, iOS, Web, Python)
Beginner friendliness	Moderate	Very high
Hardware support	Multi-band (firmware default 869.525 MHz EU; US builds ~910/915 MHz); runs many of the same LoRa boards as Meshtastic	Broad (many boards/frequencies)
Room servers	First-class firmware role (alongside Repeater and Sensor)	No direct equivalent role; nearest analogs are the Store & Forward module, channels, and MQTT bridging
Community size	Smaller, more technical	Much larger

Repeater Density and Coverage Calculations

How many repeaters do you need, and where should they go? This page provides practical calculation methods for MeshCore network coverage planning.

Link Budget Basics

The maximum range between two MeshCore nodes depends on the link budget:

```
Link Budget = TX Power + TX Antenna Gain + RX Antenna Gain - Feedline Loss - Path Loss
```

Example (typical repeater setup):

TX Power: +22 dBm (158 mW - the SX1262 PA maximum on typical LoRa boards, NOT the legal limit. FCC 47 CFR 15.247 permits up to 1 W / +30 dBm conducted in 902-928 MHz, subject to antenna-gain/EIRP rules.)

TX Antenna: +5 dBi (fiberglass omni)

RX Antenna: +5 dBi (fiberglass omni)

Feedline Loss: -1 dB each end = -2 dB total

Path Loss at 5 km free space: ~105.6 dB at 915 MHz

(FSPL = $32.45 + 20 \cdot \log_{10}(d_{\text{km}}) + 20 \cdot \log_{10}(f_{\text{MHz}})$)

Receiver Sensitivity (SX1262, SF9 / BW125 kHz): ~-129 dBm

(per Semtech SX1262 datasheet; sensitivity varies with bandwidth - e.g. ~-126 dBm at SF9 / BW250 kHz)

Available fade margin:

$(22 + 5 + 5 - 2) - 105.6 - (-129) = \sim 53$ dB fade margin

Real-world adjustment (buildings, terrain, foliage): -10 to -20 dB

(an empirical clutter/excess-loss rule of thumb; see models such as Hata/COST-231 or ITU-R P.1546)

Net fade margin: ~33-43 dB - solid link

Compliance note: Under 47 CFR 15.247, conducted output power in the 902-928 MHz band is capped at 1 W (+30 dBm), referenced to an antenna of 6 dBi or less. Antennas above 6 dBi require a dB-for-dB reduction in conducted power, yielding a derived EIRP ceiling of 36 dBm (4 W). There is **no** point-to-point antenna-gain exemption at 915 MHz - that applies only to the 2.4 and 5.8 GHz bands. The +22 dBm above is the SX1262's hardware ceiling, well within these limits; see the Antennas & RF / FCC compliance pages before increasing power or antenna gain.

Terrain Effects on Range

Free-space calculations assume line of sight. The figures below are best-case line-of-sight estimates assuming default SF/BW and ~5 dBi omni antennas. Treat them as approximate upper bounds, not planning targets - actual ranges vary widely with spreading factor, bandwidth, antenna height, power, and clutter. For emergency-grade coverage, derate heavily and confirm with field tests. Real-world path loss modifiers:

Environment	Typical Range (equal-height nodes)	Range (one node elevated 30m)
Flat open terrain	3-8 km	10-20 km
Suburban (low buildings)	1-3 km	5-10 km
Dense urban (high-rise)	0.3-1 km	2-5 km
Forest/jungle	0.5-2 km	2-5 km
Mountainous (valley-to-peak)	Variable	20-50 km (ridge-to-ridge)

The 20-50 km ridge-to-ridge figure is achievable only with full line of sight, adequate Fresnel-zone clearance, and typically *directional* antennas - not the omni setup assumed elsewhere on this page. These are approximate field-reported results, not guaranteed coverage.

Coverage Area Calculation

For a given expected range R, a single omnidirectional repeater covers approximately:

$$\text{Coverage area} = \pi * R^2$$

At R = 3 km: ~28 km² (~11 sq miles)

At R = 5 km: ~78 km² (~30 sq miles)

At R = 10 km: ~314 km² (~121 sq miles)

These are theoretical maximums. As a planning rule of thumb, actual coverage is typically only 50-70% of the theoretical circle due to terrain, buildings, and RF absorption - and real coverage is usually governed by client antenna height (handheld, low) rather than the repeater's radius.

Repeater Density Guidelines

The figures below are rough starting points for a first pass only, derived from the (unsourced, optimistic) range table above - not sourced guidelines. They assume near-best-case line of sight and ignore that most clients are handheld and low to the ground, so they tend to under-build coverage. Deploy conservatively and infill from measured RSSI/SNR data. For a network where most clients are within 1 hop of a repeater:

- **Urban dense (Manhattan, downtown Chicago):** 1 repeater per 0.5-1 km² (500m radius)
- **Suburban:** 1 repeater per 3-8 km² (1-1.5 km radius)
- **Rural flat terrain:** 1 repeater per 20-50 km² (2.5-4 km radius)
- **Rural with elevation advantages:** 1 repeater per 50-200 km² (4-8 km radius)

These are starting points biased optimistic. After initial deployment, use the actual RSSI/SNR data from your node database to identify coverage holes and place additional repeaters strategically.

Path Hop Analysis

In MeshCore, messages travel via discovered paths. The path length (hop count) determines:

- **Latency:** roughly 100-500ms per hop in normal conditions (a rough range - actual airtime depends on SF/BW/payload, and MeshCore's configurable txdelay also affects per-hop latency)
- **Reliability:** Each hop multiplies failure probability. If per-hop reliability were 95%, a 5-hop path delivers ~77% of the time ($0.95^5 = 0.774$) - but 95% per hop is optimistic. At a more realistic 80% per hop, a 5-hop path drops to ~33% (0.80^5). Per-hop reliability in real deployments varies widely with SNR, contention, and duty cycle, so do not assume 95%. For mission-critical messaging, design for 1-3 hops and verify delivery rates empirically.

Target: most clients should reach their typical destination (a room server, gateway, or key peer) within 3 hops. 5+ hops indicates a coverage gap that a new repeater could address. This 3-hop target is a reliability guideline, not a protocol limit - MeshCore supports up to 64 hops in firmware. Increasing the hop limit lets packets travel farther but does not improve per-hop reliability, so deep paths remain unreliable for time-critical traffic.

MeshCore Repeater Hardware Builds

Budget MeshCore Repeater: Under \$60

Budget MeshCore Repeater: Under \$60

Not every deployment calls for a weatherproof solar installation. For indoor sites - offices, community centers, apartment building hallways, or any location with reliable mains power - a minimal MeshCore repeater built around the RAK WisBlock platform delivers excellent performance at a fraction of the cost of a full outdoor build.

Bill of Materials

Prices below are approximate and as of early 2026; they vary by retailer, variant, and import tariffs, so treat the total as an estimate rather than a fixed figure. Check the linked vendors for current pricing.

Component	Purpose	Approx. Cost (early 2026)
RAK4631 WisBlock Core	nRF52840 + SX1262 LoRa SoC	~\$18-24 (store.rakwireless.com)
RAK19007 WisBlock Base Board (2nd Gen, USB-C)	USB-C power, slot carrier	~\$9.99 MSRP (~\$15 via US distributor)
915 MHz 5 dBi Fiberglass Antenna	Omni coverage	~\$15-25 depending on vendor
u.FL (IPEX)-to-N Pigtail (15 cm)	Antenna connection	~\$5
Total		~\$50-70

Assembly Walkthrough

Start by seating the RAK4631 into Slot A of the RAK19007 base board. The module locks with a satisfying click; verify it is fully seated and the gold contacts are flush. The RAK4631 LoRa antenna port is a u.FL/IPEX (MHF) press-fit connector, not a threaded SMA jack - so use a u.FL(IPEX)-to-N

pigtail. Snap the u.FL end straight down onto the port with gentle finger pressure until it clicks; do NOT twist or screw it, and never use pliers. Thread the N-type end through your chosen mounting point (a simple shelf bracket works well indoors) and attach the 5 dBi fiberglass antenna. Position the antenna vertically for best omni radiation.

Connect a standard USB-C cable to any 5 V / 1 A USB adapter or powered hub. The RAK19007 includes onboard power regulation; no additional circuitry is required for indoor mains-powered operation.

Firmware Flashing (REPEATER Variant)

1. Flash the RAK4631 using the MeshCore Web Flasher at flasher.meshcore.io, which serves prebuilt firmware releases (the underlying binaries are published on the official MeshCore GitHub releases).
2. Select the `REPEATER` build variant - not CLIENT or ROOM_SERVER.
3. Put the RAK4631 into bootloader mode by double-tapping the reset button; the board appears as a USB mass-storage drive (the volume name varies by bootloader, e.g. **RAK4631** or a BOOT-style name).
4. Drag the `.uf2` firmware file onto the drive. The board reboots automatically.
5. Confirm operation by connecting via the MeshCore companion app and verifying the device advertises as a repeater.

Configuration Notes

For a purely indoor repeater with AC power, no power-management tuning is required. Leave TX power at the firmware default (about 22 dBm - the SX1262 hardware maximum, well under the 30 dBm / 1 W conducted limit of 47 CFR 15.247). Your effective radiated power is TX power plus antenna gain minus cable loss; with the 5 dBi antenna here (under the 6 dBi threshold, so no power reduction is required) that is roughly 27 dBm EIRP. If you later fit a higher-gain antenna (above 6 dBi), 15.247(b)(4) requires you to reduce conducted power dB-for-dB above 6 dBi to stay compliant. Set a meaningful node name that identifies the location (e.g., `BLDG-A-3F-RPT`) so network operators can read topology at a glance.

Expected Performance

These are rough, order-of-magnitude estimates that vary dramatically with building materials, height, spreading factor, and antenna - not guarantees, so field-test in your specific building. With a 5 dBi antenna at mid-floor height, this build typically reaches roughly 300-600 m in urban environments with mixed building penetration. Clear line-of-sight across open office floors can extend to around 1-2 km. The raw LoRa PHY data rate is identical regardless of firmware role (it is set by SF/BW/CR); the REPEATER role changes forwarding and airtime behavior, not the per-packet bitrate.

Best Use Cases

- Indoor floor-by-floor mesh coverage in multi-story buildings
- Gap-fill repeaters at sites that already have AC power
- Rapid deployment for events or temporary activations
- Lab/test environments for firmware development

For outdoor, weatherproof, or off-grid deployments, see the *Pro MeshCore Solar Repeater* page in this chapter.

Pro MeshCore Solar Repeater: Complete Build

Pro MeshCore Solar Repeater: Complete Build

This guide covers a fully self-contained, weatherproof, solar-powered MeshCore repeater intended for rooftops, hilltops, and any site without mains power. Budget roughly \$230 in parts (re-priced against current vendor listings as of June 2026, after correcting the mis-specified antenna and enclosure below) and two to three hours of build time. Prices are approximate and vary by vendor and stock; confirm each line item before ordering.

Bill of Materials

Component	Purpose	Approx. Cost (as of Jun 2026)
RAK4631 WisBlock Core (nRF52840 + SX1262) + RAK19007 Base	Main radio/MCU stack	\$37 (RAK19007 base ~\$9.99; verify at the RAKwireless store)
915 MHz (902-928 MHz) Fiberglass Omni, ~5.8-8 dBi, N-type (e.g. a Rokland / Data-Alliance 900 MHz fiberglass collinear)	High-gain outdoor antenna	\$45-60
u.FL (IPEX) to N-type Pigtail, ~1 m, LMR-195/LMR-240	Low-loss feedline (u.FL end snaps onto the RAK4631 LoRa port; N-type goes to the antenna)	\$12
Genuine IP66/IP67 Polycarbonate Outdoor Enclosure (e.g. Hammond 1554/1555 series polycarbonate)	Weather protection	\$25-35
10 W Monocrystalline Solar Panel	Energy harvest	\$25
10 Ah LiFePO4 Battery Pack (12 V) with BMS that includes a low-temperature charge cutoff	Overnight/cloudy buffer	\$45
Victron SmartSolar MPPT 75/10	Charge controller	\$35

Component	Purpose	Approx. Cost (as of Jun 2026)
Inline fuse + holder for the battery positive lead (sized per the Victron table, typically the controller's rated amps)	Over-current protection	\$5
Total		~\$230 (varies by vendor)

Cite a specific SKU and dated price for each line item before ordering. The previously listed "Taoglas FXP73" is a 2.4 GHz ~2.5 dBi u.FL flex-PCB Wi-Fi antenna and does NOT belong in a 915 MHz build - it has been removed. The previously listed "Hammond 1591XXFLBK" is an ABS box rated only to ~IP54 (indoor/general use), not an IP67 polycarbonate outdoor enclosure - it has been replaced with a genuine IP66/IP67-rated polycarbonate enclosure above.

Enclosure Preparation

Begin with the genuine IP66/IP67 polycarbonate enclosure. Drill two cable-gland holes on the bottom face: one for the antenna pigtail and one for the solar/battery DC wiring. Match each gland size to the hole size and cable OD - for example, a 16 mm hole takes a PG-9/PG-11 gland; verify the gland's cable clamping range covers the pigtail's OD and the DC wiring bundle so the seal actually grips and weatherproofs. Tighten the locknuts. Place a 10 g silica gel desiccant packet inside the enclosure before sealing to prevent condensation. Mount the RAK19007 base board to the enclosure floor using four M3 x 6 mm nylon standoffs - never mount directly to the enclosure plastic, as flexing can crack solder joints.

PCB Mounting and Anti-Static Precautions

Before handling the RAK4631, ground yourself with a proper ESD wrist strap connected to a known ground. (Touching the enclosure mounting hardware only helps if that hardware is actually bonded to ground, which it usually is not - a wrist strap to a known ground is preferred.) Seat the RAK4631 into the RAK19007. The RAK4631 LoRa antenna port is a u.FL/IPEX snap-fit connector, not a threaded SMA - you cannot "screw on" an SMA jumper. Snap the u.FL end of the u.FL-to-N pigtail onto the RAK4631 LoRa port, then route the N-type end out through the cable gland (u.FL/IPEX end inside, N-type outside) and hand-tighten the gland before final enclosure assembly. Leave 5 cm of slack inside the enclosure to prevent cable strain on the delicate u.FL/IPEX connector.

Solar Wiring

Wire in the Victron-recommended sequence: connect the 10 Ah LiFePO4 battery to the BATTERY terminals of the SmartSolar MPPT 75/10 **first**, then connect the solar array to the PV terminals; disconnect in reverse order (PV first, then battery). Install an inline fuse in the battery positive lead, close to the battery terminal (sized per Victron's table, typically the controller's rated amps);

respect polarity on all connections. **Important:** the SmartSolar MPPT 75/10 has no LOAD output. Do not wire the node to a "LOAD" terminal - it does not exist on this controller. Instead, power the node from the battery: connect a DC-DC 12 V-to-5 V buck converter (set to a stable 5.0-5.1 V, rated above the node's peak TX current) to the battery terminals through a fused line or a fused distribution block, then feed its 5 V output to the RAK19007 USB-C input. Confirm the RAK19007 USB-C accepts the injected 5 V without USB negotiation. Choose wire gauge by actual current and run length: for this few-amp system over runs under ~1 m, 18-20 AWG for the battery and panel and 20-22 AWG for the load are acceptable - but check ampacity and voltage drop for your specific build. Secure all wires with zip ties to the enclosure interior.

Cold-Weather Charging Safety

Never charge a LiFePO4 (or any lithium) battery below 0 °C (32 °F). Charging below freezing causes permanent lithium plating and a real fire/safety hazard. Use a LiFePO4 pack whose BMS includes a low-temperature charge cutoff (it blocks charging below 0 °C), or add an external low-temp charge disconnect in the charge path. For cold-climate or winter sites, verify the BMS low-temp cutoff before deploying. Do not rely on any "tolerates temperature extremes" marketing - that does not mean it is safe to charge when frozen.

Antenna Installation

Rooftop safety first: before any work at height, use fall protection, never work on a wet or icy roof, and keep the mast and antenna at least their full length plus 3 m clear of any overhead power line - contact is fatal. For a fixed outdoor antenna, bond and ground the mast per NEC Article 810. Then mount the 915 MHz N-type fiberglass omni to a standard 1-inch pipe mast using a U-bolt bracket, positioned at least 1 m above the roofline for minimum ground-plane interference. Connect the N-type end of the u.FL-to-N pigtail to the antenna's N connector. Apply self-amalgamating tape over the N-type connector and at least 50 mm up the cable for weatherproofing.

Firmware and Configuration

Flash the RAK4631 with the MeshCore `REPEATER` firmware as described in the Budget Build page. After first boot, configure a fixed GPS position if known with `set lat` and `set lon` (improves network topology display). Enable the low-power sleep mode appropriate for solar-only nodes to extend overnight battery life. Set TX power with `set tx 22` (the serial CLI command; valid range is 1-22 dBm - note the reflashing page elsewhere that states 27 dBm is wrong, as 27 is out of range for the SX1262). Before settling on 22 dBm, check EIRP compliance: $EIRP = TX\ power + antenna\ gain - feedline\ loss$. Under 47 CFR 15.247(b)(4), antennas above 6 dBi require dB-for-dB conducted-power reduction below the 1 W (30 dBm) conducted limit. With an ~8 dBi antenna at 22 dBm the result is roughly 30 dBm EIRP, within the 36 dBm / 4 W EIRP ceiling; if you fit a higher-gain antenna, reduce TX power accordingly so you stay within the FCC Part 15.247 limits for 902-928

MHz. Label the node with its site name (the MeshCore name limit is 32 bytes, or 24 if a location is set) for operator reference.

Pre-Deployment Testing

Before mounting on-site, bench-test the complete assembly: cover the solar panel, run overnight to confirm the battery sustains operation for at least 12 hours, then expose to sunlight and verify MPPT charging resumes. Confirm the node is visible in the MeshCore companion app from at least 500 m away in open terrain before final installation.

MeshCore Network Expansion Strategies

When to Add a Repeater vs. When to Move One

When to Add a Repeater vs. When to Move One

Every mesh network operator eventually faces two related but distinct decisions: should you spend money on new hardware, or should you reallocate what you already have? This page gives you a structured framework for making that call.

Signs You Need a New Repeater

Adding hardware is justified when a gap in coverage is genuine and cannot be resolved by repositioning existing nodes. Look for these indicators:

- **Coverage gap identified by hop count:** If you can see, in the app's path/hop display, that users in an area consistently route through four or more hops - or experience frequent message failures - that location likely needs a node. MeshCore supports up to 64 hops in firmware, so there is no hard 3-hop protocol limit; however, delivery reliability declines with each hop because per-hop success probabilities multiply, and latency rises as paths grow. As a community rule of thumb (not a protocol spec), many operators treat roughly three hops as the point beyond which time-critical delivery becomes unreliable. Validate the real threshold for your network with measured delivery rates rather than hop count alone.
- **New neighborhood with active users:** A cluster of users has joined from an area that was previously uninhabited or where no one had a radio. Coverage there is zero - there is no node to reposition.
- **Edge-of-coverage user:** A new member joins from a location that can hear the mesh but with marginal signal. Rather than asking them to upgrade antenna hardware, a well-placed intermediate repeater improves reliability for all nearby users simultaneously.

Signs You Should Move an Existing Repeater

Repositioning is often higher-impact per dollar than buying new hardware. Consider moving a repeater when:

- **Low-traffic repeater competing for airtime:** A node that forwards very few messages per hour - check logs or room server statistics - may be co-located with a better-positioned node. Redundant coverage in the same area wastes airtime that could instead serve a gap elsewhere. Note that when two radios are co-located, receiver desensitization (front-end overload) occurs regardless of how far apart their channels are; it is mitigated primarily by physical separation, shielding, and filtering. Channel/frequency separation only addresses adjacent-channel interference, not strong-signal desense.
- **Changed environment:** A building constructed after your original deployment may now block the path your repeater was designed to bridge. Treat your network map as a living document; re-audit coverage whenever significant physical changes occur nearby.
- **Better site identified:** A rooftop access agreement, a new community partner with a tower, or simply a higher hill - if a superior site becomes available, the math usually favors moving over adding.

Cost-Benefit Framework

Before committing to either action, estimate the impact. As a rough rule of thumb (an editorial heuristic, not a fixed rule), a new repeater - roughly \$60-\$220, matching this library's Budget (~\$57) and Pro (~\$217) builds - is easy to justify when it brings a meaningful cluster of users into reliable coverage, on the order of 10-15 or more active users. But for emergency communications, even a single critical location - a responder staging area, a shelter - can justify a repeater regardless of headcount. Repositioning an existing repeater costs only your time (a few hours) and the risk of temporarily losing coverage during the move - usually 30-90 minutes. If repositioning achieves 80% of the benefit of a new node, move first and buy later.

One practical heuristic: if the candidate site for repositioning serves both the existing coverage area AND the gap, move the node. If the candidate site would leave the existing area uncovered, buy a second node to fill the gap and keep the original in place.

Documentation Practice

Record every deployment decision in a simple network log: node ID, site, date placed, reason for placement, date and reason for any relocation. This history becomes invaluable when diagnosing problems or onboarding new network operators who were not present for the original decisions.

Linking Isolated Mesh Islands

Linking Isolated Mesh Islands

As independent community mesh networks grow, they sometimes develop in parallel - two neighborhoods, two towns, or two emergency response zones that each have healthy internal mesh coverage but no connection between them. When those communities have reason to communicate or coordinate, linking the islands becomes a priority. This page covers the main technical approaches and when each is appropriate.

Option 1: Long-Range Backbone Link (Yagi-to-Yagi)

A directional point-to-point RF link between two high sites can bridge roughly 15-50 km, but only under favorable conditions: full line-of-sight, a slow spreading factor, adequate Fresnel-zone clearance, and antennas mounted high enough to clear the earth's curvature. Treat the upper end of that range as a best case, not a guarantee. Each end requires a high-gain Yagi or panel antenna (12-17 dBi is typical for LoRa backbone links), a clear line-of-sight path with adequate Fresnel zone clearance, and a dedicated MeshCore node in REPEATER mode pointed at the far end. Note that any antenna above 6 dBi triggers a dB-for-dB conducted-power reduction under 47 CFR 15.247(b)(4): at 12-17 dBi you must reduce TX power 6-11 dB below 1 W accordingly, and there is no fixed point-to-point antenna-gain exemption at 902-928 MHz (that exemption applies only to the 2.4 and 5.8 GHz bands). The maximum legal EIRP remains 36 dBm (4 W). This approach is the lowest-latency and most resilient option when geography cooperates.

Technical requirements: Calculate path loss using a link budget tool before committing to hardware. At 915 MHz, a 15 dBi antenna exceeds the 6 dBi baseline, so 15.247(b)(4) requires reducing conducted power by 9 dB - capping legal conducted output near **21 dBm** (not the SX1262's 22 dBm maximum), which keeps EIRP at or below the 36 dBm (4 W) ceiling. With that configuration the link budget can support links on the order of 40 km, but reaching that distance requires full line-of-sight with both antennas mounted high enough to clear the earth's curvature and at least 60% of the first Fresnel zone - not merely flat terrain at low height. Hills, trees, and buildings reduce range significantly. Keep at least 60% of the first Fresnel zone clear of obstructions; the first-Fresnel-zone radius (in metres) is $F1 = 17.32 \times \sqrt{(d1 \cdot d2 / (d \cdot f))}$, with $d1$, $d2$, and d in km and f in GHz, and you want obstacle clearance of at least $0.6 \times F1$. Verify any distance estimate with a link-budget tool and a field test.

When it makes sense: Two networks that share emergency response responsibility - adjacent fire districts, overlapping amateur radio emergency service areas - benefit most from a persistent RF backbone that works without Internet infrastructure.

Option 2: Internet / VPN Tunnel Bridge

When the two networks are too far apart for a clean RF line-of-sight path, or terrain blocks a backbone link, you can bridge them over the public Internet instead. A small computer (or a node with a network connection) at each island relays mesh traffic across an IP tunnel - for example an MQTT broker or a VPN/encrypted tunnel between the two sites - so packets that originate on one mesh are re-injected into the other. The RF networks stay independent; only a logical bridge ties them together.

Technical requirements: Each end needs reliable Internet access (wired, cellular, or fixed wireless) and a device to run the bridge software. Secure the tunnel (VPN or authenticated/encrypted MQTT) so the link is not open to the Internet at large, and be aware that this approach depends entirely on Internet uptime - it fails exactly when the grid or backhaul does, which matters for emergency use.

When it makes sense: Two communities separated by distance or terrain that rules out an RF backbone, where both sites already have dependable Internet and the convenience of interconnection outweighs the loss of infrastructure independence.

Option 3: Dual-Radio Bridge Node

A single physical site - ideally at high elevation between the two networks - hosts two LoRa radios, each tuned to a different mesh channel. The bridge node forwards traffic between channels, effectively stitching the two meshes together. This requires a custom firmware build or a lightweight software bridge running on an attached microcontroller or single-board computer.

Technical requirements: The bridge site must have RF visibility into both networks. Two co-located radios on the same band will desensitize each other from raw transmitter power coupling - and crucially, this front-end overload happens regardless of frequency separation. A few hundred kHz of channel offset (e.g. 500 kHz) only reduces adjacent-channel interference; it does not by itself prevent a nearby transmitter from desensitizing the co-located receiver. A practical co-site build therefore needs as much frequency separation as the band allows **plus** physical antenna separation, cross-polarization, shielding, and/or band-pass or cavity filtering, and you should avoid simultaneous TX/RX where possible. Power requirements are roughly double those of a single-radio node.

When it makes sense: Two networks that share the same general region but grew independently on different channel plans. The dual-radio bridge allows both communities to keep their existing channel configurations while gaining interconnection.

Choosing an Approach

Pick the option that matches your geography, your tolerance for Internet dependence, and the channel plans already in use:

- **Choose the RF backbone link (Option 1)** when the two sites have - or can build - clear line-of-sight from elevated locations and you need an infrastructure-independent link that keeps working when the Internet and grid are down. This is the best choice for emergency-response use, at the cost of antenna height, aiming, and FCC-compliant power planning.
- **Choose the Internet / VPN bridge (Option 2)** when distance or terrain rules out an RF path and both sites have dependable Internet. It is the easiest to deploy and maintain, but it depends entirely on backhaul uptime, so it is the weakest choice for grid-down scenarios.
- **Choose the dual-radio bridge node (Option 3)** when two nearby networks run different channel plans and you want to interconnect them without forcing either community to reconfigure. Budget for the co-site isolation engineering (filtering and physical separation) and the roughly doubled power draw.

Many real deployments combine approaches - for example, an RF backbone as the primary path with an Internet tunnel as a fallback. Whatever you choose, validate the link with a field test before relying on it.