

MeshCore Network Design

- [Planning a MeshCore Community Network](#)
- [MeshCore vs Meshtastic: Choosing for Your Community](#)
- [Repeater Density and Coverage Calculations](#)

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)
App ecosystem	Smaller	Larger (Android, iOS, Web, Python)
Beginner friendliness	Moderate	Very high

Factor	MeshCore	Meshtastic
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.