

Node Placement Strategy

Principles and methods for selecting repeater sites that maximise coverage and network reliability.

- [Repeater Placement Principles](#)
- [Coverage Radius Estimation by Terrain Type](#)
- [The Repeater Grid Approach for Urban Coverage](#)
- [Designing for Redundancy](#)

Repeater Placement Principles

The Three Rules of Repeater Placement

Every successful LoRa mesh deployment rests on three placement rules. Violate any one of them and the network will under-perform regardless of hardware quality, antenna gain, or software tuning.

Rule 1 - Height Above Terrain

LoRa operates in UHF/SHF spectrum (433 MHz, 868 MHz, 915 MHz) where propagation is quasi-optical. The radio horizon - not the visual horizon - sets the hard ceiling on range. Antenna elevation is the single most powerful lever available to a network planner.

Doubling antenna height does *not* double range; it multiplies it by roughly $\sqrt{2}$ ($\sim 1.41\times$). A node at 2 m AGL (above ground level) has a radio horizon of about 5 km on flat earth. Raise it to 20 m and the horizon extends to ~ 16 km. Raise it to 50 m and you reach ~ 25 km. Every metre of height is worth more at low elevations than at high ones - the return diminishes as you climb.

Practical target: For a community repeater, aim for a minimum of 10 - 15 m AGL. A rooftop mount, water tower, or hilltop installation should be the first option explored.

Rule 2 - Line-of-Sight to the Coverage Area

A repeater cannot serve an area it cannot "see." Because LoRa uses relatively low frequencies compared to Wi-Fi, it can diffract around moderate terrain features - but penetration through significant ridgelines or dense urban canyons is limited and unreliable. The Fresnel Zone (a football-shaped region around the direct path) must be at least 60% clear for reliable propagation; 100% clearance of the first Fresnel zone is the gold standard.

At 915 MHz the first Fresnel zone radius at the midpoint of a 10 km path is approximately:

$$\begin{aligned} r_1 &= 17.3 \times \sqrt{d / (4f)} \\ &= 17.3 \times \sqrt{10 / (4 \times 0.915)} \\ &\approx 17.3 \times \sqrt{2.73} \\ &\approx 28.6 \text{ m} \end{aligned}$$

Any obstacle taller than ~29 m at the midpoint of that path will impair the link. Use terrain analysis tools (HeyWhatsThat, SPLAT!, CloudRF, or the terrain profile view in Google Earth) to verify clearance *before* climbing a tower.

Rule 3 - Overlapping Coverage with Adjacent Repeaters

A mesh node at the edge of one repeater's coverage zone must be within range of at least one other repeater. This overlap provides two things: redundancy if one repeater fails, and smooth routing as a mobile node moves between coverage zones. Target 20 - 30% geographic overlap between adjacent repeater footprints. Less than 10% creates dead corridors; more than 50% is wasteful without adding reliability.

The Minimum Viable Repeater Principle

One well-placed hilltop or rooftop repeater at 30 m AGL can cover 300 - 800 km² of flat to rolling terrain. Ten ground-level nodes scattered across the same area will collectively cover far less, generate more channel congestion, and create a fragile, multi-hop routing mess.

This principle is counterintuitive for operators coming from Wi-Fi, where more access points always help. In LoRa mesh, channel time is shared by all nodes. Each additional ground-level node adds noise and contention while contributing minimal coverage.

“ **Rule of thumb:** Before deploying any new ground-level node, ask whether a single elevated repeater could solve the coverage problem instead. In almost every case, it can.

Dead Zone Identification Methods

Dead zones are coverage gaps where a mesh client cannot hear any repeater. Identifying them systematically before deployment prevents costly re-work later.

Method 1 - Terrain Masking Analysis

Load a digital elevation model (DEM) of your target area into a tool such as SPLAT!, Radio Mobile, or CloudRF. Run a viewshed analysis from each proposed repeater site. Areas outside all viewsheds are your primary dead zones. This takes minutes and costs nothing.

Method 2 - Wardrive / Walk Survey

Deploy a mobile node running MeshMapper or a GPS-tagged Meshtastic client. Drive or walk the target area while logging signal reports (RSSI and SNR) from known repeaters. Overlay the log on a map. Holes in the RSSI heat-map are dead zones. This method catches dead zones that terrain analysis misses (e.g., urban canyons, underground parking, indoor spaces).

Method 3 - Predictive Modelling with Terrain Obstruction Factor

After estimating coverage radius by terrain type (see the Coverage Radius Estimation page), draw coverage circles around each repeater. Grid cells covered by zero circles are dead zones. Adjust for known terrain features (a 50 m ridge between repeater and target cell reduces effective radius by roughly 30 - 50%).

Method 4 - Community Spot Reporting

For an existing network, a simple form where users report "no connectivity" locations aggregated on a shared map is surprisingly effective. Each report is a data point; a cluster of reports identifies a dead zone. The disadvantage is latency - it only works after deployment.

Using Terrain Analysis to Find Optimal Sites Before Going Outside

The cheapest field test is the one you never have to make. Modern terrain analysis tools allow a planner to shortlist candidate sites from a desk and eliminate poor candidates before anyone puts boots on the ground.

1. **Download a DEM.** USGS 1/3 arc-second DEMs (10 m resolution) are freely available for the contiguous USA via The National Map. Import into QGIS, SPLAT!, or similar.
2. **Identify candidate high points.** Use the DEM hillshade layer to find hilltops, ridgelines, and tall structures within your target area. Water towers and tall buildings are often not in the DEM - add them manually.
3. **Run viewsheds.** For each candidate, run a viewshed at your estimated antenna height (e.g., 15 m AGL). Note the percentage of target area covered.
4. **Rank candidates.** Sort by coverage area. A single site covering >80% of target area is a strong anchor candidate.
5. **Check access and permission.** The best RF site is worthless if you cannot install equipment there. Check ownership, zoning, and access before committing to a site.

6. **Field-verify top candidates.** Visit only the top 2 - 3 sites. Confirm elevation, access, power availability, and backhaul options (if applicable).

This workflow typically reduces field visits by 70 - 80% compared to ad-hoc site hunting and produces better placements because decisions are based on quantitative terrain data rather than intuition.

Coverage Radius Estimation by Terrain Type

The Radio Horizon Formula

The theoretical radio horizon for a single antenna at height h metres above a smooth spherical earth is:

$$d \text{ (km)} = 3.57 \times \sqrt{h_m}$$

For a link between two antennas at heights h_1 and h_2 the total radio horizon is:

$$d_{\text{total}} \text{ (km)} = 3.57 \times (\sqrt{h_1} + \sqrt{h_2})$$

This is the *maximum possible* range on a flat, unobstructed earth. Real terrain, vegetation, and buildings reduce this significantly. The formula provides the ceiling; terrain obstruction factors bring it down to a realistic estimate.

Terrain Obstruction Factors

Apply these multipliers to the radio-horizon distance to get a realistic coverage radius:

Environment Type	Obstruction Factor	Effective Coverage (%)	Notes
Flat open (farmland, desert, water)	0.85 - 1.00	85 - 100%	Near-theoretical range; ground reflections can help at low angles
Gentle rolling terrain	0.65 - 0.80	65 - 80%	Moderate ridge shadowing; elevated repeaters mitigate well
Suburban (low-rise, gardens)	0.50 - 0.65	50 - 65%	Houses and trees add 5 - 15 dB of excess loss
Dense forest / jungle	0.35 - 0.55	35 - 55%	Vegetation loss: 0.2 - 0.4 dB/m through canopy at 915 MHz

Environment Type	Obstruction Factor	Effective Coverage (%)	Notes
Urban (mid-rise, 3 - 8 floors)	0.30 - 0.50	30 - 50%	Building diffraction dominates; rooftop-to-rooftop paths much better
Dense urban (high-rise, canyons)	0.15 - 0.30	15 - 30%	Multipath + shadowing severe; per-block coverage planning needed
Mountainous / complex terrain	0.20 - 0.60	20 - 60%	Highly variable; valleys may have near-zero coverage from a single site

Worked Examples

Example 1 - 30 m Tower on Flat Land

A community group installs a repeater at the top of a 30 m self-supporting tower in flat agricultural land. The mobile clients they serve have antennas at 1.5 m AGL.

Radio horizon (tower): $3.57 \times \sqrt{30} = 3.57 \times 5.48 = 19.6$ km

Radio horizon (client): $3.57 \times \sqrt{1.5} = 3.57 \times 1.22 = 4.4$ km

Total radio horizon: $19.6 + 4.4 = 24.0$ km

Obstruction factor (flat open): 0.90

Realistic coverage radius: $24.0 \times 0.90 \approx 21.6$ km

Coverage area: $\pi \times 21.6^2 \approx 1,466$ km²

With a single well-sited tower, this repeater comfortably covers a circle of about 1,500 km² - roughly the size of a mid-sized US county.

Example 2 - 10 m Mast in Suburbs

A volunteer mounts a repeater on a 10 m mast attached to their house in a typical American suburb (one-storey houses, trees). Mobile clients at 1.5 m.

Radio horizon (mast): $3.57 \times \sqrt{10} = 3.57 \times 3.16 = 11.3$ km

Radio horizon (client): $3.57 \times \sqrt{1.5} = 3.57 \times 1.22 = 4.4$ km

Total radio horizon: $11.3 + 4.4 = 15.7$ km

Obstruction factor (suburban): 0.55

Realistic coverage radius: $15.7 \times 0.55 \approx 8.6$ km

Coverage area: $\pi \times 8.6^2 \approx 232$ km²

This is a solid community repeater covering roughly the footprint of a small city. However, pockets of shadow behind larger buildings will exist. A wardriving survey is recommended after installation to confirm actual coverage.

Example 3 - Rooftop in Dense Urban

A repeater is placed on the flat roof of a 7-storey (21 m) apartment building in a dense city. Surrounding buildings average 6 storeys (18 m). Mobile clients at street level (1.5 m).

Radio horizon (roof): $3.57 \times \sqrt{21} = 3.57 \times 4.58 = 16.4$ km

Radio horizon (client): $3.57 \times \sqrt{1.5} = 3.57 \times 1.22 = 4.4$ km

Total radio horizon: $16.4 + 4.4 = 20.8$ km

Obstruction factor (dense urban): 0.25

Realistic coverage radius: $20.8 \times 0.25 \approx 5.2$ km

Coverage area: $\pi \times 5.2^2 \approx 85$ km²

Only 25% of the theoretical range is realised because the repeater barely clears the surrounding roofline. Even 3 additional storeys of height (to 30 m) would raise the obstruction factor to ~ 0.35 and expand coverage to ~ 120 km². In dense urban networks, rooftop height relative to surroundings matters enormously.

Coverage Overlap: The 20 - 30% Rule

When planning adjacent repeaters, their coverage footprints should overlap by 20 - 30% of the coverage radius. This ensures:

- No gap corridor between repeaters where nodes lose connectivity
- Sufficient signal margin at the cell edge for reliable forwarding (not just barely detectable signals)
- Redundancy: a node in the overlap zone can reach two repeaters

If repeater A has a coverage radius of 8 km and repeater B also has 8 km, place the two sites no more than 13 - 14 km apart (leaving 2 - 3 km of overlap on each side). Placing them 16 km apart creates a 0 - 2 km gap corridor in the middle where nodes may lose one or both repeaters.

Link Budget Margin: Target >10 dB

Coverage radius estimates are not hard boundaries - they define the distance at which the link margin drops to zero. In practice, a link operating at exactly the sensitivity floor is unreliable. Fading, multipath, vegetation sway, and atmospheric changes will cause it to fail intermittently.

Target at least 10 dB of margin at the cell edge for reliable operation. This means planning for a coverage radius at which the received signal is 10 dB above the receiver sensitivity floor.

For Meshtastic on SF12 / BW125, receiver sensitivity is approximately -137 dBm. A link budget target of -127 dBm at the coverage boundary gives 10 dB margin. Links measured below -127 dBm at normal operating distances should be treated as marginal and either reinforced with a relay node or deprioritised until a better repeater site is available.

Quick field check: If a node reports $RSSI < -125$ dBm or $SNR < -10$ dB when communicating with its nearest repeater, that link is at or below the 10 dB margin boundary. Plan to add a relay or move the repeater closer.

The Repeater Grid Approach for Urban Coverage

Why a Grid Approach?

Ad-hoc repeater placement - putting nodes wherever a willing host can be found - produces uneven coverage with clusters of overlapping repeaters in some areas and dead zones in others. A systematic grid approach starts from coverage requirements and works backward to site requirements, ensuring uniform coverage and efficient use of channel capacity.

Target Repeater Density by Area Type

Area Type	Target Repeater Spacing	Approx. Density	Rationale
Urban core (high-rise, dense)	1.0 - 2.0 km	1 per 1 - 3 km ²	Heavy building obstruction limits each repeater to a small footprint
Suburban (low-rise, residential)	3.0 - 5.0 km	1 per 7 - 20 km ²	Moderate obstruction; rooftop repeaters can reach 3 - 8 km reliably
Exurban / light industrial	5.0 - 8.0 km	1 per 20 - 50 km ²	Low building density; tall structures (grain elevators, water towers) excellent
Rural (farmland, grassland)	8.0 - 15.0 km	1 per 50 - 175 km ²	Near line-of-sight; hilltop or tower sites dominate
Wilderness / remote	15.0 - 30.0 km+	1 per 175 - 700 km ²	Solar-powered mountaintop repeaters; realistic in open/semi-open terrain only

These spacings assume repeaters are elevated (10 - 30 m AGL for urban, 30+ m for rural) and achieve 10 dB of link margin at the target spacing. Reduce spacing by 30 - 40% if using low ground-mounted installations.

Step 1 - Identify Anchor Sites

Anchor sites are high, prominent structures that provide disproportionately large coverage footprints. Finding and securing anchor sites is the most important work in urban coverage planning. Common anchor site types:

- **Water towers:** Typically 25 - 50 m AGL, publicly owned, often with existing antenna mounts. Utility companies are sometimes receptive to emergency communications partnerships. Ideal anchor nodes.
- **Tall commercial buildings (10+ storeys):** Roof access is harder to obtain but a single mid-rise rooftop repeater in a dense urban core can match 5 - 10 lower installations. Target telecommunications companies, building management firms, or property owners sympathetic to community projects.
- **Hilltops and ridgelines:** In cities built on rolling terrain (Pittsburgh, San Francisco, Seattle), hilltops within or adjacent to urban areas are the highest-value sites. Even a 5 - 10 m height advantage over surrounding terrain dramatically extends range.
- **Radio / TV transmission towers:** Licensed broadcasters and tower companies sometimes offer shared mounting space. The collocation fee may be justified by the coverage gained.
- **Church steeples and clock towers:** Often the tallest structures in older residential neighborhoods. Many faith communities are receptive to community emergency communications projects.

Map all potential anchor sites in a GIS tool. Run viewshed analyses from each. Rank by coverage area. Secure the top 3 - 5 sites before planning secondary fill repeaters.

Step 2 - Fill Planning

Once anchor sites are installed and their actual coverage verified (wardriving survey or community signal reports), overlay the confirmed coverage zones on your planning map. Grid cells with no coverage, or with RSSI below -120 dBm, are fill targets.

Fill repeaters do not need to be as elevated as anchor sites - they only need to bridge a specific gap. A 5 m residential rooftop install may be sufficient to cover a neighborhood dead zone if it is positioned on the line of sight between two anchor sites. Prioritise fill sites that:

1. Cover the largest dead zone area with the smallest number of new nodes
2. Can hear at least two anchor-tier repeaters (for redundancy)
3. Are accessible for maintenance

Step 3 - Documenting Coverage Gaps

Maintain a living coverage map updated after every wardriving run or community signal report. A shared GIS layer (e.g., a Google MyMaps or an QGIS project shared via cloud storage) accessible to all network administrators is the best tool for this. For each coverage gap, record:

- Geographic boundary of the gap (polygon on the map)
- Estimated area (km²) and population affected
- Identified candidate fill site(s), if any
- Status: *open / candidate identified / site secured / fill node deployed*
- Last confirmed date of the gap (wardriving date or community report date)

A structured gap log prevents the common failure mode of deploying fill nodes in areas that are already well-covered while neglecting persistent dead zones that lack a vocal advocate.

Worked Example: Planning a Mid-Sized City Grid

Target: A city of 80,000 people covering approximately 120 km², mostly suburban with a 2 km² dense downtown core.

1. **Downtown core (2 km²):** Target 1 - 2 km spacing → need 1 - 2 anchor repeaters. Identified a 12-storey bank building and a water tower 1.3 km apart. Both anchor sites secured. Downtown coverage achieved with 2 nodes.
2. **Suburban ring (118 km²):** Target 4 km spacing → grid of approximately 7 - 8 repeaters needed. Identified 6 water towers and 2 church steeples evenly distributed across the suburban ring. All 8 sites secured. Average spacing: 3.9 km.
3. **Total anchor infrastructure: 10 nodes** for 120 km² = 1 repeater per 12 km².
4. **Post-wardriving fill:** Survey revealed 3 dead zones in valley neighborhoods. 3 fill repeaters added on rooftops in those valleys. Total: 13 nodes for complete city coverage.

Compare this to an ad-hoc approach: a typical volunteer-driven deployment in a city this size might have 40 - 80 ground-level nodes with significant overlap in connected areas and persistent dead zones in underserved neighborhoods. The grid approach delivers better coverage with far fewer nodes and less channel congestion.

Designing for Redundancy

Why Single-Path Mesh Is Fragile

A tree-topology mesh - where each node has exactly one path back to the network core - is the natural shape that forms when coverage is barely adequate. In a tree topology, the failure of any interior node partitions the network. Nodes "below" the failed repeater become an isolated island: they can hear each other but cannot reach the wider network.

For a community mesh serving emergency communications, this is unacceptable. The very events that trigger heavy mesh use (storms, earthquakes, infrastructure failures) are also the events most likely to take repeaters offline through power loss, physical damage, or access loss.

N+1 Redundancy

N+1 redundancy means that every coverage area is served by at least N+1 repeaters where N is the minimum needed for coverage. In practice, for a mesh network, this translates to:

“ Every node in the network should be able to reach the network core via at least two independent paths through different physical repeaters.

A "network core" is the set of nodes with Internet gateway access or the central coordination point. In a city mesh, the core might be 3 - 5 well-connected anchor repeaters.

To verify N+1 for a given node:

1. Identify all repeaters that node can directly reach ($RSSI > -120$ dBm).
2. For each of those repeaters, confirm it has at least one other path back to core.
3. If the node can only reach a single repeater, it has no redundancy. If that repeater fails, the node is isolated.

Ring Topology vs. Tree Topology

Tree Topology

Nodes connect to the nearest repeater, which connects to the nearest anchor, which connects to core. This forms a tree. Advantages: simple to plan and understand. Disadvantages: any broken branch isolates all nodes below it. Single points of failure are everywhere.

Core

```
├─ Anchor A
  │ ── Repeater A1
  │ │ ── Client nodes (isolated if A1 fails)
  │ ── Repeater A2
├─ Anchor B
├─ Repeater B1
└─ Client nodes (isolated if B1 or Anchor B fails)
```

Ring Topology

Anchor repeaters are interconnected in a ring so that each anchor has two paths back to core. Fill repeaters connect to two anchors where physically possible. This creates a lattice rather than a tree.

Core — Anchor A — Anchor B — Anchor C — Core

\ | /

Repeater Fill Repeater

(hears A (hears (hears B
and B) B and C) and C)

Ring topology requires more careful planning and more anchor sites (each anchor must be within radio range of two others), but it eliminates the single points of failure that make tree topologies fragile.

Recommendation: Design anchor-tier repeaters in a ring or lattice. Fill-tier repeaters can remain in a simplified tree to anchor, but each fill node should reach at least two anchors where terrain permits.

Identifying Single Points of Failure with Path Analysis

A single point of failure (SPOF) is any node whose failure disconnects part of the network. Identify SPOFs through path analysis:

1. **Draw the network graph.** Each repeater is a node. Each radio link between repeaters is an edge. Include only links with at least -115 dBm (15 dB margin; reliable links only).
2. **Find cut vertices.** A cut vertex is any node whose removal disconnects the graph. In graph theory this is computed with DFS (depth-first search). In practice, you can identify them visually: any node that is the sole bridge between two subgraphs is a cut vertex and therefore a SPOF.
3. **Prioritise SPOF mitigation.** For each SPOF identified, either add a redundant link (find a fill repeater position that bypasses the SPOF) or ensure the SPOF node has UPS backup power, weatherproof housing, and remote monitoring.

Testing Redundancy by Taking a Node Offline

Theoretical redundancy analysis should be validated with live tests. The procedure is simple:

1. **Notify operators.** Announce a planned maintenance window (e.g., "Node X will be taken offline for 30 minutes on Saturday 14:00 UTC for redundancy testing").
2. **Take the target node offline** by powering it down or disconnecting its antenna.
3. **Measure impact.** Using a network map (MeshMapper, Meshtastic node list, or MeshCore admin panel), observe which nodes lose connectivity. Nodes that disappear from the map are isolated - this is your actual failure impact, which may differ from the theoretical prediction.
4. **Document the partition.** Record which nodes were isolated and for how long they would be unreachable in a real failure event.
5. **Restore the node** and plan remediation for any isolated segments found.

Perform redundancy tests at least once per year, and after any significant change to the network topology (adding or removing anchor repeaters, significant coverage expansion).

Practical Redundancy Checklist

- Every anchor repeater can reach at least two other anchors directly.
- Every fill repeater can reach at least two anchors directly.
- No anchor repeater is a single point of failure for more than one fill repeater.
- All anchor repeaters have UPS or generator backup covering at least 72 hours.
- Network graph has been drawn and cut vertices identified.
- Redundancy live-test performed in the last 12 months.
- Failure impact documented: "If node X fails, Y nodes lose connectivity."