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 spectrum where propagation is quasi-optical. The band depends on region: 902 - 928 MHz in the United States, 863 - 870 MHz in the EU, and 433 MHz in some regions. In the US, use 902 - 928 MHz — 868 MHz is the EU ISM band and is not available for Part 15 LoRa here, and 433 MHz US operation is tightly restricted under Part 15. (Note: 433/868/915 MHz are all UHF; SHF begins at 3 GHz.) 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 √2 (~1.41×). Using the standard refraction-corrected (4/3-earth) radio horizon formula **d(km) ≈ 4.12 × √(h in metres)**, a node at 2 m AGL (above ground level) has a radio horizon of about 5.8 km on flat earth. Raise it to 20 m and the horizon extends to ~18 km. Raise it to 50 m and you reach ~29 km. Every metre of height is worth more at low elevations than at high ones - the return diminishes as you climb. (The purely geometric/optical horizon uses the smaller coefficient 3.57 × √h; the 4.12 figure includes normal atmospheric refraction.)

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 serves areas it can "see" far better than areas it cannot. Because LoRa uses relatively low frequencies compared to Wi-Fi, it can diffract around moderate terrain features - but non-line-of-sight coverage via diffraction is weak and unreliable, and penetration through significant ridgelines or dense urban canyons is limited. 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:

```
r₁ = 17.32 × sqrt(d / (4f))   [equivalently 8.66 × sqrt(d / f)]
 = 17.32 × sqrt(10 / (4 × 0.915))
 ≈ 17.32 × sqrt(2.73)
 ≈ 28.6 m
 (d in km, f in GHz)
```

By the 60% rule, an obstacle begins to impair the link once it intrudes more than ~40% into the first Fresnel zone - i.e., once it rises into roughly the central ~17 m of the ~29 m radius at midpoint (about 11 - 12 m of intrusion). Full clearance of all ~29 m is the gold standard. 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. As a planning rule of thumb, target roughly 20 - 30% geographic overlap between adjacent repeater footprints. These are illustrative heuristics rather than measured thresholds: too little overlap (very roughly under ~10%) tends to leave dead corridors, while very heavy overlap (over ~50%) adds little extra reliability for the extra sites.

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 ridge intruding the line-of-sight path adds diffraction loss that depends on the ITU-R P.526 diffraction parameter (the ridge's height relative to the direct line, its position along the path, and frequency), so model it per-path rather than applying a fixed percentage. A ridge near one endpoint, at the midpoint, or barely grazing the Fresnel zone all produce very different losses.

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 can substantially reduce the number of field visits required compared to ad-hoc site hunting, and produces better placements because decisions are based on quantitative terrain data rather than intuition.

Revision #3
Created 2026-05-03 05:37:31 UTC by Mesh America Admin
Updated 2026-06-09 22:31:21 UTC by Mesh America Admin