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 desensing 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.

Revision #5
Created 2026-05-03 06:58:50 UTC by Mesh America Admin
Updated 2026-06-09 01:01:24 UTC by Mesh America Admin