

MeshCore Network Topology Best Practices

MeshCore Network Topology Best Practices

Backbone vs. Client Layer

A well-designed MeshCore network is organized into two distinct layers:

- **Backbone layer:** dedicated repeaters placed on elevated sites with clear line-of-sight between them. These form the routing backbone that carries traffic across the network. They are the infrastructure - always on, high antenna, fixed location.
- **Client layer:** user devices (phones, handhelds, base stations) that connect to the nearest backbone node. In MeshCore, companion (client) nodes are endpoints, not relays - per the MeshCore FAQ, clients do not repeat traffic for other nodes.

This two-layer separation keeps the forwarding load on the backbone. Clients add no *forwarding* load, but they still transmit their own messages, advertisements, and path-discovery floods on the shared channel, all of which consume airtime. So adding clients does not add relay load to the backbone, but a very dense client population can still congest the shared channel and indirectly affect performance.

Repeater Placement Guidelines

The numbers below are rules of thumb for planning, not protocol limits - tune them to your terrain and traffic.

- Aim for **3 - 5 repeaters per coverage zone**, each with line-of-sight to at least 2 others in the backbone.
- **Avoid single points of failure** - if one repeater goes offline, the network should remain functional via alternate paths.

- Ensure **overlapping coverage** between adjacent repeaters so that clients are never more than 1 hop from the backbone.
- High sites (hilltops, building rooftops, water towers) dramatically extend backbone range - prioritize elevation over raw transmit power.

Hop Budget

MeshCore supports up to 64 hops (the protocol ceiling). As a planning rule of thumb, aim for **no message traversing more than 6 - 8 backbone hops**. Beyond this:

- Per-hop latency accumulates noticeably.
- Each additional hop adds another potential failure point.
- Route re-discovery after a link failure takes longer with more hops in the chain.

For wide-area networks that would otherwise require long hop chains, use **room servers as message hubs** rather than relying on extended peer-to-peer relay chains.

Advertisement Tuning

- **Flood advertisements** (visible network-wide) should be infrequent - **every 12 hours** is appropriate for stable infrastructure nodes. Frequent floods waste airtime and provide no benefit when the topology is static.
- **Zero-hop advertisements** (local only, for client discovery) can be more frequent - every few minutes is reasonable.
- Review your advertisement intervals if you observe unexplained airtime congestion on the channel.

Mesh Segmentation for Large Networks

In a very large network (50+ repeaters), avoid trying to relay everything peer-to-peer across the entire mesh. Instead:

- Use **room servers as message hubs** for cross-region delivery. Room servers provide message storage and delivery confirmation.
- Segment the mesh into regional clusters, each with its own backbone, connected via room servers at the regional boundaries.
- This reduces the hop count needed for cross-region delivery and localizes the impact of any regional topology change.

Monitoring Topology Health

The MeshCore app includes a **network map feature** that shows which repeaters a node can see and the routes between them. Use this to:

- Verify backbone connections are healthy after deployment.
- Identify repeaters that have lost contact with their neighbors (indicates a failure or coverage gap).
- Confirm that new repeaters have been discovered and integrated into the routing fabric.
- Check hop counts for key routes and identify bottleneck nodes carrying disproportionate traffic.

Revision #3

Created 2026-05-03 04:19:16 UTC by Mesh America Admin

Updated 2026-06-09 14:28:30 UTC by Mesh America Admin