

Understanding LoRa Mesh Networking

Core concepts: what mesh networking is, how LoRa works, platform comparisons, and US frequency regulations.

- [What Is a Mesh Network?](#)
- [LoRa Technology Explained](#)
- [LoRa vs LoRaWAN vs Meshtastic vs MeshCore](#)
- [The 915 MHz ISM Band](#)

What Is a Mesh Network?

If you have ever used Wi-Fi at home, you are already familiar with the most common type of wireless network: the **star topology**. Every device in your house - your phone, your laptop, your smart TV - talks to one central access point (your router), and the router connects everything to the internet. It is simple, it works well indoors, and it is fine as long as that one router keeps working.

But what happens if the router goes offline? Every device loses connectivity at the same instant. The whole network collapses around a single point of failure. Now imagine you are in a disaster zone, or deep in a national forest, or at a community event with no cellular service. A star network is useless. There is no router to plug in, and even if there were, its failure would silence everyone.

A **mesh network** solves this by making every single device both a *user* and a *router*. Instead of everyone talking through one hub, every node talks to the nodes near it, and those nodes pass the message onward, hop by hop, until it reaches its destination. Remove any one node - or even several - and traffic finds another path around the gap, provided another path exists - in a sparse mesh, losing a single well-placed relay can still partition the network. There is no central authority and no central server that can take the whole network down; however, a sparse mesh can still depend on individual relay nodes.

How a Message Travels in a Mesh

Picture five hikers spread across a mountain trail, each carrying a LoRa radio. Alice is at the trailhead; Echo is at the summit. Between them, out of direct radio range of each other, are Bob, Carol, and Dave.

1. Alice wants to send Echo a message: "*Summit still clear?*"
2. Alice's radio broadcasts the packet. Bob and Carol can both hear it.
3. Bob rebroadcasts it. Carol rebroadcasts it. (Duplicate suppression logic prevents infinite loops - each packet carries a unique ID that nodes remember and discard if seen again.)
4. Dave hears one of those rebroadcasts and passes it along.
5. Echo receives the packet, having never been able to hear Alice's original transmission directly.
6. Echo's reply travels back by whatever path is available at that moment.

This is called **multi-hop packet forwarding**. Each intermediate node is sometimes called a *relay* or a *repeater*. The exact path a message takes is not fixed - it depends on which nodes are currently powered on and within radio range of each other. The network continuously *self-heals*: if Dave's battery dies mid-hike, messages between Carol and Echo simply reroute through whoever is still on the air.

Self-Healing Topology

The term **self-healing** means the network does not need a human administrator to reconfigure routes when a node disappears. Routing happens automatically. The flooding-based protocols this wiki covers (Meshtastic, MeshCore) do not maintain explicit routing tables - nodes simply rebroadcast and the surviving nodes carry traffic along whatever paths remain. When a node disappears, the mesh stops depending on it without any user intervention; how quickly traffic reroutes depends on the protocol and how the network is configured.

This is why mesh networks are so valuable for emergency communications, outdoor adventures, festivals, and anywhere you cannot rely on existing infrastructure.

Star Topology vs. Mesh Topology at a Glance

Feature	Star (Wi-Fi, cellular)	Mesh (LoRa mesh)
Central hub required?	Yes - everything depends on it	No - every node is equal
Single point of failure?	Yes	No central one (but a sparse mesh can still hinge on one key relay)
Range extension	Buy another access point, manually configure	Add any node; routing is automatic
Works off-grid?	Only if you have power and hardware	Yes - battery-powered nodes, no internet needed
Works during infrastructure failure?	No	Yes - if nodes have independent power and remain in radio range
Setup complexity	Plug in the router and done	Flash firmware, configure - slightly more involved

Why Mesh Is Ideal for Off-Grid Communications

In situations where commercial infrastructure is unavailable or unreliable - wildfires, earthquakes, backcountry recreation, sailing, amateur radio events - a mesh network can be a useful supplemental option for short- to medium-range text communications, alongside (not instead of)

proven tools like two-way radios or satellite messengers with SOS for life-safety use. Here is why:

- **No subscription fees.** There is no carrier, no monthly bill, no account needed.
- **No internet connection required.** The mesh works entirely peer-to-peer. The nodes themselves are the network.
- **Incrementally deployable.** Even two nodes form a functioning (if very small) mesh. Every node you add extends the network further.
- **Low power.** LoRa radios can run for days or weeks on a small battery. Nodes can be solar-powered for permanent outdoor deployment.
- **Long range.** A single LoRa hop can span several kilometers in open terrain. With a few well-placed elevated relay nodes a multi-hop mesh can cover a county - though each added hop consumes shared channel capacity, and Meshtastic's default hop limit (3, configurable up to 7) must be considered.

The next page explains the radio technology that makes all of this possible at such low power and over such long distances: LoRa.

LoRa Technology Explained

LoRa stands for **Long Range**. It is a wireless radio modulation technique invented by a French startup called Cycleo and acquired by Semtech in 2012. LoRa chips appear in millions of devices worldwide today, from smart utility meters to wildlife trackers to community mesh nodes. To understand why LoRa is so effective for mesh networking, you need to understand the core technique it uses: **Chirp Spread Spectrum**.

What Is Chirp Spread Spectrum?

Most radio systems transmit data on one fixed frequency. A chirp spread spectrum (CSS) system does something different: it sweeps the signal continuously across a range of frequencies, rising or falling in pitch over time (like a bird's chirp - hence the name). The receiver looks for this specific sweep pattern rather than listening at one spot.

Why does this matter? Several reasons:

- **Noise rejection.** Man-made interference tends to be concentrated at specific frequencies, while a CSS signal is spread across the whole channel. Because the signal is spread across a wide band, any noise at one frequency only affects a small fraction of the signal. The receiver can mathematically reconstruct the original data even when much of the band is noisy.
- **Multipath resistance.** Radio signals bounce off buildings, hills, and trees. These echoes arrive at the receiver slightly delayed and can cancel out the original signal. CSS is much more resilient to this effect than narrowband modulation.
- **Low detectability.** A spread-spectrum signal looks almost like background noise to anyone not specifically looking for it, making it robust in RF-noisy environments.

Processing Gain and Extraordinary Sensitivity

The key metric for a radio receiver is its **sensitivity** - the minimum signal power it can decode. A conventional FM radio might have a sensitivity of around - 100 dBm (decibel-milliwatts). LoRa achieves sensitivities around - **134 to - 137 dBm** at the bandwidths used for mesh (BW125-250 kHz, SF12), and as low as ~- 148 dBm only at very narrow bandwidth (7.8 kHz) not used by Meshtastic/MeshCore. That is roughly 50 dB better than FM, which means LoRa can receive signals

tens of thousands of times weaker.

This extraordinary sensitivity comes from **processing gain**: the mathematical process of correlating the received sweep against an expected template. The more time and bandwidth the receiver spends correlating, the more gain it recovers. LoRa lets you tune this trade-off with a parameter called the **Spreading Factor**.

What Spreading Factor Actually Does

The Spreading Factor (SF) is a number from 5 to 12 on modern LoRa chips (SX126x); older SX127x chips are used with SF7-12. It controls how many chips are used to encode each symbol (and each symbol carries SF bits).

- **Low SF (e.g., SF7):** fewer chips per symbol → faster data rate → shorter range → lower battery use per packet. Used when nodes are close together.
- **High SF (e.g., SF12):** more chips per symbol → very slow data rate → very long range → more battery use per packet. Used when you need maximum range.

Think of it like this: shouting a word once at normal speed (SF7) versus repeating every syllable ten times very slowly (SF12). The slowly repeated version can be understood even with a lot of background noise, but it takes much longer to say.

A common community mesh preset uses **SF10** or **SF11**, which balances range and throughput for typical text messaging workloads. Meshtastic's default "LongFast" channel uses SF11 on 915 MHz.

The Range / Speed / Power Triangle

In radio, you almost never get something for nothing. LoRa is no exception. The three variables - range, data rate, and power consumption - are always in tension:

- To go farther (more range), you must slow down (lower data rate) or transmit louder (more power).
- To go faster (higher data rate), you must accept shorter range or higher power.
- To save power, you must accept either shorter range or slower speed.

LoRa's key achievement is that it pushes this triangle to extremes: it can achieve very long range at very low power, but only by accepting a very slow data rate. A LoRa packet might carry 50 - 250 bytes of useful data. The raw over-the-air data rate at SF12 is about 250 bits per second. That is slower than a 1990s dial-up modem. But for text messages, GPS coordinates, and sensor readings, it is entirely sufficient - and few technologies match that range at that power level in unlicensed spectrum.

LoRa vs. FSK vs. GFSK

To appreciate LoRa, compare it with the modulation techniques used by competing low-power radios:

Modulation	How it works	Typical sensitivity	Typical range (open field)	Used in
FSK (Frequency Shift Keying)	Switches between two fixed frequencies for 0 and 1	- 112 dBm	~1 - 2 km	Many 433/915 MHz modules, older APRS
GFSK (Gaussian FSK)	FSK with a Gaussian filter to reduce bandwidth	- 105 to - 115 dBm	~1 - 3 km	Bluetooth, ANT, older Nordic proprietary radios
LoRa (CSS)	Frequency sweep across entire channel bandwidth	- 137 to - 148 dBm	5 - 15+ km	Meshtastic, MeshCore, LoRaWAN

The ~30 - 40 dB sensitivity advantage is large: in ideal free space a 30-40 dB advantage corresponds to 30-100x the distance of FSK at the same transmit power. Over real terrain the practical gain is smaller (several times the range) but still dramatic. This is why LoRa took over the low-power, long-range wireless market.

Link Budget: The Simple Version

A **link budget** is just an accounting exercise: you add up all the gains and losses between transmitter and receiver and check whether the signal is still strong enough at the end. Here is a simplified example:

```
Transmit power: +30 dBm (1 watt)
Transmit antenna gain: + 2 dBi
Path loss (10 km): -120 dB (including ~8 dB excess/ground loss over the ~112 dB free-space value)
Receive antenna gain: + 2 dBi
-----
Received signal: -86 dBm
```

LoRa sensitivity (SF11, 250 kHz as used by LongFast): -131 dBm
Link margin: +45 dB (plenty of headroom)

A positive link margin means the link can work; in practice you want 10-20 dB of spare margin for fading. A larger margin means you have room to lose (walls, foliage, less-than-perfect antenna placement). The extraordinary sensitivity of LoRa gives you a huge margin even for difficult paths through vegetation or inside buildings.

LoRa Is Not LoRaWAN

One important clarification: LoRa is just the physical-layer radio modulation. It says nothing about how multiple devices share the channel, how data is addressed, or how the network is organized. **LoRaWAN** is one specific network protocol built on top of LoRa - but it is not the only one. Meshtastic and MeshCore are entirely different protocols, also built on LoRa. The next page compares all four of these in detail.

LoRa vs LoRaWAN vs Meshtastic vs MeshCore

One of the most common sources of confusion when starting out is the alphabet soup of acronyms: LoRa, LoRaWAN, Meshtastic, MeshCore. People use them interchangeably, but they are four very different things operating at four different layers. This page explains each one clearly and shows you when each is the right choice.

The One-Line Summary of Each

- **LoRa** - a radio modulation technique (the physics of the signal)
- **LoRaWAN** - a centralized IoT network protocol built on top of LoRa
- **Meshtastic** - an open-source, peer-to-peer mesh protocol built on top of LoRa
- **MeshCore** - another open-source, peer-to-peer mesh protocol built on top of LoRa

LoRa is the foundation. Everything else is a different architecture built on top of that foundation.

Comparison Table

	LoRa	LoRaWAN	Meshtastic	MeshCore
What is it?	Radio modulation (physical layer only)	Network protocol + cloud infrastructure	Open-source mesh firmware + apps	Open-source mesh firmware + apps
Network topology	N/A - just the radio signal	Star-of-stars: end nodes → gateways → cloud server	True peer-to-peer mesh	True peer-to-peer mesh
Requires gateway?	N/A	Yes - a LoRaWAN gateway is mandatory	No	No
Requires internet?	N/A	Yes - gateway must reach a network server	No	No

	LoRa	LoRaWAN	Meshtastic	MeshCore
Who manages it?	Semtech (chip maker)	LoRa Alliance (standards body) + network operators	Open-source community	Open-source community (created by Scott Powell; clients by Liam Cottle)
Primary use case	Any LoRa-based application	IoT sensor reporting (meters, trackers, sensors)	Community mesh, off-grid text + GPS	Community mesh, off-grid text + GPS
Typical data payload	Anything LoRa can carry (up to ~250 bytes)	Small sensor readings (10 - 50 bytes)	Text messages, GPS, telemetry	Text messages, GPS, telemetry
Works off-grid?	N/A	No in practice - requires a gateway plus a network server (almost always internet-hosted)	Yes - fully off-grid	Yes - fully off-grid
Encryption	N/A - you implement it	AES-128 end-to-end (mandated by spec)	AES-256 channel encryption	Elliptic curve, per-message encryption
App ecosystem	None (raw hardware)	Industry dashboards (The Things Network, etc.)	Android, iOS, web client	Android, iOS, web client
Beginner-friendly?	Hardware only - requires firmware development	Moderate - requires gateway setup and TTN account	Yes - free flasher tool, polished apps	Yes - free flasher tool, active community

LoRa: The Physical Foundation

LoRa is implemented inside a Semtech radio chip (such as the SX1276 or SX1262). The chip generates and receives the chirp-spread-spectrum signal. By itself, a LoRa chip does nothing useful - it is like having a blank radio transmitter with no protocol to tell it what to say or how to talk with others. You always need a higher-level protocol on top of it.

LoRaWAN: Built for IoT, Not for People

LoRaWAN was designed to connect large numbers of low-power *sensors* to the internet. The architecture is deliberately star-shaped: sensors (called "end nodes") are not allowed to talk to each other directly. They can only talk upward to a gateway, and the gateway forwards data to a cloud-based "network server." LoRaWAN is excellent for reading water meters, tracking shipping

containers, or monitoring soil moisture across a farm - all cases where data flows in one direction toward a central server.

LoRaWAN is *not* designed for:

- Direct person-to-person messaging
- Off-grid use (without a powered gateway with internet access)
- Community resilience networks (a device loses connectivity when no reachable gateway or the network server is down - the network is only as resilient as its infrastructure)

Meshtastic: Open-Source Mesh for the People

Meshtastic is a free, open-source project that turns cheap LoRa hardware into a fully off-grid text and GPS mesh network. Developed since 2020 and now maintained by a large open-source community, it is the most widely deployed LoRa mesh platform in the world. Key characteristics:

- Nodes communicate directly with each other using a flooding-based mesh routing algorithm called **Managed Flood**.
- Users control the network via polished Android and iOS apps connected to the node over Bluetooth or Wi-Fi.
- There are no gateways required. Any node automatically acts as a relay for other nodes.
- Large community, extensive documentation, active development.
- Best choice if you want to join the largest existing mesh community or buy from the widest hardware selection.

MeshCore: A Newer Alternative with Different Trade-offs

MeshCore is a newer open-source mesh platform with a slightly different design philosophy. Rather than simple message flooding, MeshCore uses a more structured routing approach that can reduce channel congestion on larger, busier networks. Key characteristics:

- Roles are explicit: nodes are configured as *clients*, *repeaters*, or *room servers* (group chat hubs). This clarity reduces unnecessary retransmissions.
- Strong encryption using elliptic-curve cryptography.
- Also has Android, iOS, and web clients.

- Growing community, particularly in areas building more deliberately structured networks.
- Good choice if you want more control over how traffic is routed or if you are building infrastructure (repeaters, room servers) rather than just carrying a node.

Which One Should You Use?

- **You want to buy sensors and report data to the cloud:** Use LoRaWAN with The Things Network.
- **You want to join the biggest community mesh and get started as fast as possible:** Use Meshtastic.
- **You want a more structured network with explicit roles and you do not mind a slightly steeper learning curve:** Use MeshCore.
- **You want to build your own custom firmware from scratch:** Use raw LoRa hardware and write your own protocol on top of it.

For most beginners, Meshtastic is the right first step. MeshCore is an excellent second platform to explore once you have a feel for how mesh networking works in practice. Both are covered in detail throughout this wiki.

The 915 MHz ISM Band

Nearly all LoRa mesh devices sold for North America operate in the **915 MHz ISM band** (902 - 928 MHz). (A few 433 MHz LoRa devices also exist and are usable here, but they are uncommon.) Understanding what that means - and what the rules are - will help you choose the right hardware, set the right channels, and avoid interference with your neighbors.

What Is the ISM Band?

ISM stands for **Industrial, Scientific, and Medical**. These bands are designated internationally (ITU Radio Regulations) for industrial, scientific and medical RF applications. They are not, strictly, "set aside for unlicensed communications" - rather, in the US the FCC additionally allows unlicensed communications devices (like LoRa) to share them under Part 15 rules, on a non-interference, secondary basis. (Note: 902-928 MHz is an ISM band only in ITU Region 2, the Americas.) The trade-off is that these bands are open to many users, and everyone has to play nicely together by following power limits and other technical rules.

In the United States, the band used by LoRa spans **902 to 928 MHz** (commonly referred to as the "900 MHz band" or "915 MHz band"). It is regulated by the FCC under **Part 15** of the Code of Federal Regulations.

FCC Part 15 Power Limits

The FCC imposes strict limits on how much power you can transmit in this band:

- **1 watt (30 dBm) conducted power** - this is the maximum power at the antenna connector of the radio.
- **4 watts (36 dBm) EIRP (Equivalent Isotropically Radiated Power)** - this is the derived ceiling that results from the gain-reduction rule, not a separate flat allowance. The 1 W conducted limit assumes an antenna gain of up to 6 dBi. If your antenna gain *exceeds* 6 dBi, you must reduce conducted power dB-for-dB by the amount the gain exceeds 6 dBi (per 47 CFR 15.247(b)(4)). An antenna of 6 dBi or less requires no power reduction.

What does this mean in practice? Most LoRa modules transmit at 20 - 27 dBm (0.1 - 0.5 watts). A typical 3 dBi gain antenna is perfectly legal at full transmit power. A 10 dBi antenna (4 dB above 6 dBi) requires reducing conducted power by 4 dB, to 26 dBm - the reduction is keyed to antenna gain, not to any EIRP arithmetic. (For fixed point-to-point links the reduction is more lenient, 1 dB

per 3 dB of gain above 6 dBi.) Almost no consumer LoRa hardware comes close to these limits, so for most users, this is a non-issue.

No License Required (With Caveats)

Because LoRa operates under Part 15, **you do not need an amateur radio license or any other license** to operate a Meshtastic or MeshCore node in the United States. This makes community mesh networks accessible to everyone, not just licensed ham radio operators.

However, Part 15 devices must accept all interference and must not cause harmful interference to any authorized radio service - whether in-band (including primary and government users of 902-928 MHz) or in adjacent bands. In practice, the 900 MHz band is busy with cordless phones, baby monitors, some Wi-Fi equipment, and other ISM devices - but LoRa's spread-spectrum nature makes it naturally robust against narrowband interference from these sources.

Duty Cycle Considerations

In the United States, digitally modulated systems in the 902 - 928 MHz band - the category that covers LoRa mesh devices - have *no* duty-cycle limit. (Frequency-hopping systems, by contrast, do have per-channel dwell-time limits under 15.247(a)(1)(i).) LoRa itself does not frequency-hop (it stays on one channel per packet), and the Part 15 rules permit continuous operation as long as power limits are respected.

That said, good network citizenship means keeping your transmit duty cycle low. If every node on a channel is transmitting constantly, collisions will degrade performance for everyone. Meshtastic and MeshCore both implement built-in duty cycle management and back-off algorithms to prevent nodes from saturating the channel.

Channel Selection and Frequency Hopping

Within the 902 - 928 MHz band, LoRa devices can use many different center frequencies (channels). Meshtastic does not use a single fixed frequency for its default **LongFast** preset; instead it computes the exact channel frequency from the selected region, preset (modem config), and channel name, dividing the band into numbered slots. For the US region, the LongFast primary

slot lands near 906 - 907 MHz, but the precise value is derived by the firmware rather than being a hard-coded number - see the Meshtastic frequency-slot documentation. MeshCore's USA/Canada preset uses similar frequencies.

Frequency hopping (rapidly jumping between channels) is permitted and used by some competing technologies (like the older FHSS radios), but it is not required for LoRa and is not used by Meshtastic or MeshCore in their standard modes. Instead, they use a fixed channel, relying on LoRa's spread-spectrum nature to handle interference.

Channel selection matters when:

- You want to create a private channel separate from the public mesh.
- You want to avoid interference from other mesh users or industrial equipment.
- You are deploying multiple networks in the same area and need them to coexist.

What About Europe and Other Regions?

The 915 MHz band is specific to the Americas. In Europe, LoRa community mesh devices typically use the **868 MHz ISM band** (863 - 870 MHz), regulated by the ETSI under different rules including duty-cycle limits on many sub-bands. Other regions have their own band plans:

Region	LoRa Band	Frequency Range	Key Rule
USA / Canada	915 MHz	902 - 928 MHz	1W conducted / 4W EIRP, no duty cycle limit
Europe / UK	868 MHz	863 - 870 MHz	25 mW ERP and 0.1-1% duty in most sub-bands; 500 mW ERP / 10% duty in 869.4-869.65 MHz (the Meshtastic EU_868 default)
Australia / NZ	915 MHz	915 - 928 MHz	1W EIRP
Asia (varies widely)	Varies by country	Varies	e.g. China 470-510 MHz, India 865-867 MHz, Japan/Korea ~920-923 MHz, Southeast Asia 920-925 MHz (AS923)

This is critical: a European 868 MHz LoRa device will not work on a US 915 MHz mesh, and vice versa. Always check that the hardware you buy is rated for your region's frequency band before purchasing. Hardware sold in the US is built for the 915 MHz band, but Meshtastic firmware ships with the region UNSET and will not transmit until you select your region (US) in the app on first setup. If you are buying from overseas vendors, double-check the product listing.

Some newer hardware (such as devices using the Semtech SX1262 chip) can be configured in software to cover both 868 MHz and 915 MHz, as the chip supports a wide frequency range. However, the antenna is typically tuned for one band or the other, so even if the chip can transmit on the wrong frequency, performance will be degraded.