RG3 METER COMPANY

Mesh or Mess?

Why Mesh Network Architecture Struggles in Battery-Powered Water Metering

By Lee Gregory · RG3 Meter Company

In utility metering since 1998.

A Technical White Paper · Updated and expanded edition, 2026 · Originally published 2015

Executive Summary

Mesh networking is a sound architecture in the environment it was designed for: line-powered electric metering. Ported into water metering, the same architecture inherits a constraint it was never built to tolerate — a finite battery — and the very features that make mesh attractive on paper become liabilities in the field.

The core findings of this paper:

- **The selling points are real, but conditional.** Reduced infrastructure, repeater-based coverage, and self-healing redundancy all work when nodes have unlimited power. Water endpoints do not.
- **Self-healing is paid for in battery.** Every reroute, retry, and re-convergence is a transmission, and frequent transmission is precisely what kills a water endpoint's battery.
- **Throughput collapses with hops.** Half-duplex relaying loses roughly 50% of effective bandwidth per hop, which is why mesh vendors cap hop counts.
- **The receiving environment degrades as the network grows.** More relays mean more transmissions competing in the same RF space, compounded by multipath fading.
- **The meter pit is the hardest RF environment in the industry,** and a low-power mesh endpoint must escape it and relay for its neighbors on the same starved power budget.
- **"No single point of failure" inverts in practice,** producing concentrated, hard-to-diagnose failure clusters that are invisible from the head-end.

This is not an argument that mesh "doesn't work." It is an argument that mesh transfers the cost of coverage onto the one resource a water endpoint cannot replenish, and that a verified line-of-sight architecture concentrates that cost where the physics is hardest instead of spreading it across an endpoint population that has no power to spare.

1. How Mesh Is Supposed to Work

In a mesh network, each Meter Interface Unit (MIU) is also a repeater. Rather than transmitting directly to a collector, an MIU hands its data packet to a neighboring MIU, which relays it to the next, daisy-chaining until the packet reaches a collector, receiver, or gateway. By definition, a mesh network reconfigures around broken or blocked paths by retransmitting node-to-node until the destination is reached, so it can remain operational when an individual node or link fails. [1]

The advertised advantages follow directly from that design:

- **Less fixed infrastructure** — endpoints carry the network instead of towers.
- **Lower per-hop power** — each transmission only has to reach the next neighbor.
- **Line of sight matters less** — every MIU is a repeater, so packets route around obstructions.
- **No single point of failure** — data can hop around a failed node.

Each of these is true in a network that can afford to transmit freely. That qualifier is the entire subject of this paper.

2. The Power Problem: Why Electric Mesh Does Not Translate to Water

Mesh networking matured in electric metering, where every meter is wired to mains power. A line-powered node can transmit, listen, relay, and re-converge as often as the protocol demands without consequence, because its energy budget is effectively unlimited.

Water and gas meters have no such luxury. They sit in pits, vaults, basements, and wells where mains power is generally unavailable, so they run on long-life lithium batteries expected to last well over a decade. [2] A limited power supply directly limits how much information an endpoint can afford to communicate, which lowers the read-success rate that mesh depends on.

The industry has acknowledged this distinction in its own architectures. In hybrid electric-plus-water deployments, the electric meters form and carry the mesh; the water and gas meters merely report into it and do not relay other meters' data — because the frequent transmissions required to support a mesh would drain their batteries. [3] Independent technical references state the same conclusion plainly: relaying meter data requires significant power, which is a non-issue for electric AMI but a real constraint for battery-powered water endpoints. [4]

This is the foundational problem. A mesh asks every node to be a relay. A water node cannot be a full relay without spending the one resource it cannot get back.

3. Throughput and Latency: The Half-Duplex Penalty

A second problem is structural to wireless relaying and independent of battery chemistry. Most endpoints use a single, half-duplex radio: they can transmit or receive, but not both at once. To relay a packet, a node must first receive it and then retransmit it, which consumes at least twice the airtime of a direct transmission. [5]

The practical result is a well-documented rule of thumb across the wireless industry: each wireless hop reduces effective throughput by roughly 50%, and the degradation compounds with every additional hop. [6] Equipment vendors outside metering reach the same number independently and explicitly recommend limiting hop counts as a result. [6] Academic treatments express the same effect more formally — in a linear multi-hop chain, end-to-end throughput falls roughly in proportion to 1/h, where h is the number of hops. [7]

Mesh metering systems live with this by capping the number of hops a packet may take. The cap controls the throughput loss, but it also bounds how far a mesh can reach before it needs another collector — quietly eroding the "less infrastructure" selling point.

Latency rises in parallel. Each hop adds contention delay and retransmission overhead, so even where data eventually arrives, it can arrive late — a recognized reason that "near-real-time" mesh systems often deliver only previous-day interval data in practice. [8]

4. The Receiving Environment: Noise, Contention, and Multipath Fading

A mesh does not just move packets; it adds transmitters to the air. Every hop introduces another transmission into the same RF space, and as the count of simultaneous transmissions rises, the environment grows noisier and reception harder. The analogy is conversational: two people in a quiet room (a direct line-of-sight link) versus two people trying to be heard across a crowded one (a multi-hop mesh).

This is compounded by multipath fading, a propagation effect present in all RF systems but more punishing in a noisy, reflective environment. Multipath occurs when a signal reaches the receiving antenna by two or more paths — caused by reflection, refraction, diffraction, and scattering off ground, water, and structures — and the copies arrive out of phase. [9] The interfering copies produce both constructive and destructive interference and phase shifting; destructive interference causes the signal to fade, sometimes below the threshold needed to receive it at all. [9] The same physics produces the familiar "ghosting" once seen in analog television. [9]

Most unlicensed mesh metering systems operate in the 900 MHz ISM band using frequency-hopping spread spectrum. By the admission of the patent literature in this field, such systems — often engineered by communications specialists rather than meter people — can be highly complex and difficult to troubleshoot. [10] Operating in an unlicensed band also means sharing spectrum with other consumer and industrial devices, with no regulatory protection from their interference. [4]

5. "No Single Point of Failure" — Reexamined

The self-healing claim is real on paper, and it was written for line-powered networks. In a battery-powered water deployment, the mechanism that delivers resilience is the same mechanism that drains the endpoints, and the resilience degrades as a result.

Consider what a constrained endpoint actually "knows." A low-power node has minimal routing intelligence and limited memory; it keeps only a short record of the neighbors a packet has recently passed through, not a live, global picture of the network. That is enough to relay, but not enough to reason about topology. Three failure modes follow directly:

- **Routing loops.** With no global view, a packet can pass from node 1 to 2 to 3 to 4 to 5 and back to 1 — a circle that never reaches the backhaul. Mesh protocols suppress loops with hop-count limits, but the cost of suppression is identical: the packet is eventually dropped after every node in the loop has burned transmit energy carrying it nowhere.
- **The cure consumes the patient.** Each relay, retry, and re-convergence attempt is a transmission — and frequent transmission is exactly what drains water endpoints, which is why some vendors keep these meters off the relaying burden entirely. [3] A network "healing" itself by rerouting is a network spending battery it cannot replace. The endpoints can exhaust themselves trying to move a message along.
- **Orphaned, not dead.** A low-power node that loses its path must find a new one. If it lacks the transmit power to reach an alternate neighbor, it cannot re-converge at all — it simply goes silent while remaining physically functional.

From the head-end, this is where the claim inverts. A node that is genuinely dead and a node that is merely failing to complete the daisy chain look identical — both stop reporting. The troubleshooting complexity these systems are known for [10] is not incidental: when several endpoints fall silent, there is no clean way to tell which one is the actual fault, or how many faults there are. The operator is left diagnosing a silence.

And the silence concentrates. Suppose three endpoints each serve as the path to the backhaul for ten downstream meters. Lose those three relays, and thirty-three endpoints go dark at once — three faults producing thirty-three non-reporting meters across multiple collapsed branches. The architecture that promised “no single point of failure” has instead produced many concentrated points of failure, each invisible until a technician walks the route.

The honest framing: mesh does self-heal — in a network that can afford it. A water mesh cannot. Every resilience feature is paid for in battery and in operator visibility, and in a power-starved, low-intelligence endpoint population, the bill comes due as dead batteries, orphaned nodes, and unattributable outages.

6. The Meter Pit: RF Penetration and Attenuation

The line-of-sight question is often decided before the signal ever leaves the ground. A water endpoint usually transmits from inside a below-grade pit, through some combination of a lid, soil, and standing water — three of the worst dielectric environments for a low-power radio.

The metering industry documents this directly: radio signals generated from below-ground water metering pits are absorbed, refracted, and reflected, producing poor RF propagation — and when the pit has a metal lid, propagation is impacted further still. [11] Those lids are most commonly cast iron, [12] the worst case for a radio trying to escape. The governing physics:

- **Metal reflects and absorbs.** A conductive lid acts as a partial shield. At metering frequencies the effect is severe: smart meters operate roughly between 900 MHz and 2.4 GHz — wavelengths of about 12–33 cm — and even a conductive mesh with millimeter-scale openings blocks the large majority of that energy. [13] A solid cast-iron lid is not a perforated mesh; it reflects and absorbs far more, leaving only fringe coupling around the lid edges to carry the signal out.
- **Water and wet soil attenuate by conductivity.** Soil and water do not merely block a signal; they absorb it, and the loss rises with moisture and dissolved-ion content. The governing quantity is skin depth — how far the field penetrates before it decays. The dielectric properties of soil change with moisture, and increasing moisture directly reduces the skin depth of penetration. [14] Field and laboratory studies confirm that RF attenuation in soil increases with both water content and electrical conductivity. [15] This is why an endpoint that reads reliably in a dry plastic pit can disappear after a rain or in a flooded vault.
- **Higher frequency is penalized.** The unlicensed designs that lean on mesh sit at 900 MHz and above, where penetration through earth and water is poorer than at lower licensed frequencies. The endpoint pays this loss on top of the metal-lid loss and the wet-soil loss — and it pays it on a battery budget that cannot be increased.

A mesh endpoint is therefore asked to do the hardest RF job in the industry — escape a metal-capped, often-wet underground enclosure — at the lowest available power, and then relay for its neighbors on that same starved budget.

7. Mesh in the Water Market Today

Mesh is not a hypothetical architecture in water metering; it is actively marketed by established manufacturers, who describe it in their own materials.

- **Mueller Systems** built its Mi.Net Advanced Metering Infrastructure on a wireless-mesh foundation, having licensed an energy-efficient wireless-mesh patent portfolio to create what the company described at the time as a high-performance AMI mesh system. [16] The Mi.Net system operates in the unlicensed 900 MHz band. [17] Mueller has since expanded its endpoint portfolio to include LoRaWAN and cellular options alongside its 900 MHz network. [17]

- **Zenner** markets a “Mesh Fixed network system” for AMI that, in the company’s own description, locates the best available frequency, routes around obstacles, and repeats signals between units. [18]

These descriptions are presented here as published fact, not criticism. Both companies are credible manufacturers, and mesh is a legitimate engineering choice with real strengths. The point of this paper is narrower and technical: when mesh is implemented on battery-powered endpoints, the architecture’s strengths are purchased with a resource — transmit energy — that those endpoints cannot replenish.

8. The Engineering Alternative: Verified Line-of-Sight

No architecture is free. A line-of-sight or point-to-multipoint fixed network makes an opposite set of trade-offs, and honesty about them is what makes the comparison credible:

- **It requires more infrastructure** — collectors and gateways must be positioned and powered.
- **It requires more transmit power per endpoint** to reach a collector directly rather than a nearby neighbor.

In exchange, it removes the failure modes that matter most in water:

- **Endpoints do not relay,** so battery is spent only on the endpoint’s own reads — not on carrying the network.
- **There is no hop-count throughput penalty** and no multi-hop latency stacking.
- **Failures are individually visible.** A non-reporting endpoint is one endpoint, diagnosable on its own, not a silent cluster of unknown size.
- **Licensed spectrum is available.** A licensed private fixed network (for example, in the FCC Part 90 450 MHz band) operates with regulatory protection from outside interference, unlike the shared, unlicensed 900 MHz ISM band common to mesh designs. [4]

The decisive engineering principle is the same one that closes the original argument: RF is often treated as “black magic” because results can surprise, but in reality the laws of physics govern the outcome. The most durable AMR/AMI systems are those that design interference out of the network and then confirm, through verification, that each transmission was actually received — rather than assuming a packet found its own way through an unmonitored chain of hops.

Architecture comparison at a glance

Factor	Battery-Powered Mesh	Verified Line-of-Sight
Endpoint power use	Relays for neighbors; drains battery	Reads own data only
Throughput	~50% loss per hop; capped hops	No per-hop penalty
Latency	Stacks with each hop	Single link to collector
Failure visibility	Silent clusters; cause unclear	One endpoint, individually diagnosable
Spectrum	Typically unlicensed 900 MHz ISM	Licensed private band available (450 MHz Part 90)

Factor	Battery-Powered Mesh	Verified Line-of-Sight
Infrastructure	Lower fixed infrastructure	More collectors / gateways

9. Conclusion

Mesh networking earns its reputation in line-powered electric metering, where unlimited energy lets a node relay, reroute, and re-converge at will. Water metering removes that assumption. On a battery, every mesh feature becomes a withdrawal from an account that cannot be refilled: self-healing drains it, relaying drains it, retries drain it — and the operator cannot always tell the difference between a node that has run dry and a node that has simply lost its way to the backhaul.

The water meter environment is one of the most hostile in all of RF. The defensible engineering response is not to ask each undersized, underpowered endpoint to also serve as network infrastructure, but to concentrate transmit power and intelligence where the physics is hardest, operate on protected spectrum, and verify receipt of every read. That is the difference between a network that assumes it healed and one that confirms it delivered.

References

1. Mesh network self-healing definition, AMI patent literature, U.S. Patent 8,370,697.
2. AMI water meter power and battery life, BMAG Meter technical overview (long-life lithium, 6–15 year service life).
3. The Successful AMI Marriage: When Water AMR and Electric AMI Converge, WaterWorld (Elster EnergyAxis — water/gas meters report into the electric mesh but do not relay, to avoid battery drain).
4. Advanced Metering Infrastructure overview, Engineering Training Institute (relaying power burden; licensed 450–470 MHz vs. unlicensed 900 MHz).
5. Wireless mesh half-duplex throughput degradation, U.S. Patent 8,724,530 (single-radio relay requires ~2× channel time); NIST, Throughput and Delay Analysis of Half-Duplex IEEE 802.11 Mesh Networks.
6. Per-hop ~50% bandwidth loss, TP-Link mesh performance FAQ and CommScope RUCKUS community guidance (recommend limiting hop counts).
7. End-to-end throughput in multi-hop wireless networks (throughput proportional to 1/h in linear topologies), ResearchGate.
8. Constrained mesh bandwidth and previous-day interval data, Utility Dive (smart metering bandwidth limits).
9. Multipath propagation and fading, ScienceDirect (Multipath Component) and Wikipedia (Multipath propagation); reflection / refraction / diffraction / scattering, constructive and destructive interference, phase shift, and “ghosting.”
10. 900 MHz ISM frequency-hopping mesh complexity, U.S. Patent 8,279,778 (“highly complex and difficult to troubleshoot”).
11. Below-ground pit RF absorption / refraction / reflection and metal-lid impact, U.S. Patent 9,748,651.
12. Cast-iron meter pit lids, U.S. Patent 8,350,719.
13. Conductive shielding of 900 MHz–2.4 GHz signals, ShieldYourBody smart-meter shielding overview.
14. Soil moisture and skin depth of RF penetration, U.S. Patent 9,857,316.
15. RF path-loss / attenuation in soil increases with moisture and conductivity, MDPI Sensors (Comparison of RF Path Loss Models in Soil) and University of Birmingham M.Phil. thesis (Abdorahimi, 2018).
16. Mueller Systems / SIPCO energy-efficient wireless-mesh patent license for AMI, Mueller Systems press release (2011).
17. Mi.Net 900 MHz operation, Mueller Systems Mi.Net RF safety document and muellersystems.com (Mi.Net Mobile / Migratable Solution).
18. Zenner “Mesh Fixed network system” description, zennerusa.com/mesh.

Prepared by Lee Gregory, RG3 Meter Company. Technical claims in this paper are sourced to published industry, academic, vendor, and patent literature as cited above. Competitor architectures are described from the manufacturers’ own published materials.