Episode Book

EPISODE #5767

Inside Private Grids That Run Themselves

PUBLISHED	7 October 2026 at 22:39 UTC
DURATION	27:00
LISTEN	myweirdprompts.com/episode/private-microgrids-islanding-explained/
HOSTS	Corn, Herman
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.2
BOOK GENERATED	19 October 2026 at 17:39 UTC

SYNOPSIS

Most people never think about the electrical network behind a hospital, airport, or remote mine — but those private grids already do everything a utility does, just behind a fence. This episode maps the spectrum from a campus that simply buys power at one meter to a fully islandable microgrid that generates, stores, and balances its own electricity. Along the way: point of common coupling, protection coordination, anti-islanding, seamless transfers, reverse power flow lessons from Hawaii's rooftop solar saturation, and why remote mines running on diesel at three to ten times grid cost are the closest thing to true energy independence. The goal isn't fantasizing about secession from the utility — it's asking how feasible it would be for an ordinary neighborhood to adopt a model that already exists at the edge of town, humming behind a fence.

TRANSCRIPT



Corn

Daniel's got a two-parter for us this time, and it starts from a gripe.



Herman

A gripe?



Corn

Not a gripe exactly. He says we touched on community microgrids before, the neighbourhood collectively generating and storing its own power, and he wants to come back to it. But he doesn't want to start from the futuristic version. He wants to start from the boring version that already exists.



Herman

Which is most of it.



Corn

Right. He points out that a community microgrid is one example, and not even the most common one. University campuses, industrial complexes, airports, hospitals, military bases, mines, data centers, remote settlements, all of those can have substantial electrical networks under private management. Some just distribute power they buy from the public grid. Others generate their own, store it, and can disconnect entirely and run on their own.



Herman

And that's question one. What do these things actually look like in the real world, where does the public grid end, what does the private operator manage, how do generation and storage get folded in.



Corn

Then question two. How far can you push that model toward real independence? Starting from a site that only manages its own internal wires, what's the progression to a true microgrid that generates, balances, islands, and potentially runs autonomously long term. And for a residential community trying to do it, how do ownership, metering and payment work alongside the technical side.



Herman

That's a lot of question.



Corn

He also says the goal isn't to fantasise about secession from the utility. It's to look at what already exists and ask how feasible it would be for a neighbourhood to adopt the same model.



Herman

Which is the right way in, honestly, because the technology isn't speculative. It's sitting behind a fence at the edge of town, humming.



Corn

So let's start by mapping the territory. What these private grids actually are, and where the fence line sits between the public utility and the private operator.



Herman

The cleanest way to think about it is a spectrum. At one end you've got a site that takes power from the public grid at a single point and distributes it internally. That's a hospital, a data center, a shopping mall, a university. No generation of its own to speak of. It's just a customer with a really big service entrance.



Corn

One meter.



Herman

One point of interconnection, at least. The utility brings its line to a substation on the property, and from there everything on the other side of that fence belongs to the site. The transformers stepping it down, the switchgear, the medium-voltage distribution loops running between buildings, the protection relays, all of it. The utility's responsibility ends at the meter, or at the substation fence if it's a bigger service.



Corn

And that point has a name.



Herman

Point of common coupling. It's the boundary. Everything upstream is the utility's problem, everything downstream is yours. You own the failure modes.



Corn

Which is a different relationship than most people have with electricity. Most of us experience the grid as a wall socket that either works or doesn't.



Herman

And a campus experiences it as a machine with a hundred components, most of which they bought.



Corn

So at the other end of the spectrum?



Herman

A fully islandable microgrid. Own generation, own storage, own control system, and the ability to disconnect from the public grid and keep running. Not for a few seconds on a UPS, for days or indefinitely if the design allows it. That's the far end.



Corn

And the interesting stuff lives in the middle.



Herman

The middle is where almost all of it lives. Most private grids are somewhere between the two, and they don't advertise it.



Corn

Let's take the most common case first. Big campus, industrial site, takes power at one high-voltage interconnection and fans it out internally.



Herman

Take a university. A decent-sized one will have its own substation where the utility's transmission or distribution line lands, and then a medium-voltage loop, usually something like twelve and a half kilovolts, sometimes five or fifteen, running around campus in buried duct banks. Each building has its own transformer stepping down to four hundred eighty volts, and inside the building you're back to ordinary distribution.



Corn

So the campus has effectively built a miniature version of what the utility does.



Herman

That's exactly what it is. It's a small utility with a campus culture. And critically, they've taken on the operational responsibilities. They run switching operations. They have a high-voltage crew, or they contract one. They do protection coordination, which means making sure that when a fault happens, the right breaker opens and not the whole



Herman (continued)

campus.



Corn

Protection coordination being the art of not taking out three buildings because a squirrel died in one transformer.



Herman

Nobody talks about it, but yes. Selectivity. You want the device closest to the fault to trip first. Get that wrong and every small fault becomes a campus-wide outage, which is a great way to find out who your protection engineer actually was.



Corn

That's a lesson people learn once.



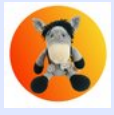
Herman

Once, expensively. And that's the part that's invisible from outside. When a campus has its own grid, it also has its own ops burden. Somebody has to be on call at two in the morning when a feeder locks out.



Corn

Now add generation.



Herman

This is where it gets more interesting. Hospitals are a good example. A hospital has its own distribution, its own transformers, but almost every hospital also has diesels for the life-safety branches, and increasingly a combined heat and power plant. CHP. You burn gas in an engine, you get electricity and you capture the waste heat for hot water and steam.



Corn

Which a hospital happens to need a lot of.



Herman

Sterilisation, laundry, heating, domestic hot water. So the heat isn't wasted, which is why CHP units make sense for hospitals and universities and industrial plants where there's a constant thermal load. A hospital running on gas with heat recovery can hit eighty percent overall efficiency, where the grid plus a gas boiler is doing maybe half of that.



Corn

And the CHP isn't just backup.



Herman

That's the part people miss. It's running all the time, in parallel with the grid, offsetting imports. It's baseload. The diesels are the backup. Two different machines with two different jobs.



Corn

So the private grid has a power plant on it.



Herman

A small one, but yes. It's synchronised to the utility, it's exporting or importing depending on the balance inside the fence, and it's providing voltage support to the campus distribution. The site is no longer purely a consumer. It's a participant.



Corn

Which changes the operational picture a lot.



Herman

A lot. Once you've got generation on your side of the fence, you've got a parallel source, you've got potential for back-feed, and you've got to think about what happens when the utility side goes away. Your generator will happily try to energise the utility line if you don't have protection against it, and that's how you electrocute a lineman on a pole.



Corn

That's anti-islanding.



Herman

Anti-islanding protection. The inverter or the generator relay has to detect grid loss and disconnect, so you don't back-feed the distribution network while somebody's out there with a bucket truck. That's the baseline requirement just to have generation at all.



Corn

So there's a minimum level of sophistication the moment you install a generator, before you get any of the good stuff.



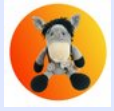
Herman

And the next step up from that is where you decide you want the site to keep running when the utility is down. That's the actual microgrid.



Corn

Airports, military bases, some data centers.



Herman

Airports are a good example. A large airport has its own distribution network, often a substantial one. Terminals, runways, taxiway lighting, fuelling systems, baggage handling. It has critical loads that cannot go dark, like air traffic control, runway lighting, fuel systems.



Corn

And it has diesels sitting on the site doing nothing most of the time.



Herman

Doing nothing, and being tested once a month, and costing money. The microgrid argument for an airport is that you use the assets you've already paid for. If you can island, you can keep the critical loads up indefinitely and the non-critical loads up as long as fuel lasts.



Corn

What's the hard part?



Herman

The transfer. Detecting that the grid has failed, opening the breaker to the utility, and transitioning to island mode without dropping your critical loads. That's seamless islanding. If it takes ten seconds, you've lost some equipment and you've had a bad day.



Corn

And if the grid comes back, you have to resynchronise.



Herman

Resynchronise, match phase, frequency, voltage, and close. That has to be done carefully or you cause a wreck on both sides. This is where the microgrid controller earns its price. It's a system that watches generation, storage, load and the utility interface in real time and balances them.



Corn

Which sounds a lot like what a utility control room does.



Herman

It is exactly that, scaled down. Which brings us to the part I find most interesting, which is that a private grid operator ends up doing the same job as a utility, with fewer people and less institutional memory.



Corn

Tell me what that job actually contains.



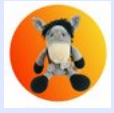
Herman

In the Hawaiian case you can see it clearly. There was a huge thread among utility people and solar installers when rooftop solar took off there, and the list of things the operator takes on the moment distributed generation gets large is long. Frequency regulation. Voltage control. Protection coordination. Managing reverse power flow, because when there's more generation on the distribution circuit than there is load, power flows the wrong way up the feeder.



Corn

Reverse power flow being the one that surprises people.



Herman

It surprises people because the grid was never designed for it. Regulators, transformer taps, protection settings, all assumed power flowed one direction, from the substation out. The moment rooftops on a circuit are exporting, some of that equipment is being asked to work backwards, and not all of it enjoys it.



Corn

Hawaii being the case where you got to see it under load.



Herman

Hawaii has enormous rooftop solar penetration, among the highest per capita on earth, and they hit limits the mainland hadn't hit yet. Circuits saturated. Utilities started requiring inverters with advanced grid support functions. They had to actively curtail. Voltage problems showed up on feeders that had too much midday export and not enough load to absorb it.



Corn

And the same physics will apply to a neighbourhood microgrid, just with fewer people around to explain it.



Herman

Fewer people and less money. That's the real asymmetry. A utility has engineers who do this for a living and a rate base to pay for it. A private grid operator has to either hire that skill or acquire it, and the cost lands on whoever owns the site.



Corn

Now let's talk about the extremes, because Daniel asked about mines and remote settlements, and those are the most interesting cases to me.



Herman

They're the closest thing that exists to a fully autonomous community microgrid. A remote mine, somewhere three hundred kilometres from anything, is not connected to a public grid in any meaningful sense. It is an island.



Corn

It's been an island the whole time.



Herman

The whole time. Historically it runs on diesel. You truck or fly fuel in, you burn it in generators, and you accept the cost. And the cost is enormous, because remote diesel generation is somewhere between three and ten times the cost of grid power depending on how remote and how bad the logistics get.



Corn

What's changing?



Herman

Solar plus storage plus, increasingly, wind. The economics started making sense in the last decade because the alternative to solar is not grid power, it's diesel, and diesel at a remote site is brutally expensive. A solar farm and a battery bank that displaces half your diesel consumption is a very different proposition when your baseline is five hundred dollars a megawatt hour and a fuel truck that takes two days to arrive.



Corn

And if the truck doesn't arrive?



Herman

You find out how committed you actually were to running the crusher. The logistics are the hidden system. It's not just generation and distribution, it's the fuel supply chain, and the site's production schedule is hostage to a road that might wash out.



Corn

So the most independent grids in the world are the ones with the fewest illusions about it.



Herman

They know what their fuel costs because somebody signs the purchase order every month. They know how many hours the generator runs. They know exactly what an outage costs in tonnes not moved.



Corn

Which is the opposite of how most of us experience the grid.



Herman

Which is why those sites are the best analogue for what a community microgrid would actually be. Not the elegant, optimised, self-balancing system. The version with a fuel delivery you have to schedule and a maintenance log you have to keep.



Corn

And the data center case, which is the other extreme.



Herman

A data center is a private grid optimised for something slightly different. Not just uptime, but power quality and continuity. They'll typically take two independent feeds from the utility, sometimes from separate substations, on separate rights of way. On-site generation, usually diesel, sized to run the whole facility. UPS systems, typically batteries, sized to bridge the gap between grid loss and generator start, which is somewhere in the ten to thirty second range.



Corn

And they test all of it constantly.



Herman

Constantly. A large data center is testing its generators, its UPS, its transfer switches on a rolling schedule. The whole point is that the day the grid fails, everything that's supposed to happen has happened a hundred times already.



Corn

That's expensive.



Herman

It's expensive and it's the cost of the service they sell. If your service level agreement says ninety nine point nine nine nine percent, you don't get to decide that redundant feeds are overkill. The customer decided that when they signed.



Corn

So we've got the whole spectrum. A campus with internal distribution only, a hospital with CHP, an airport or base that can island, a data center with triple redundancy, a mine that's completely alone.



Herman

And the common thread is that in each case somebody took on the operational responsibilities of a utility, in exchange for control, for reliability, for economics, or because they had no choice.



Corn

So that's the landscape. Now let's trace the progression from simple internal distribution to full electrical independence.



Herman

I think of it as four stages. Stage one is internal distribution only. You're a customer with your own wires. You take power at the point of common coupling, you run it around your site, and you have no generation of your own. You might have emergency backup, but it's not integrated.



Corn

No generation means no anti-islanding problem, no parallel operation, no reverse power flow.



Herman

Right, and no dependence on your own engineering. Your exposure to grid failure is total. The grid goes out, you're dark, except for whatever's on the life-safety system.



Corn

Stage two.



Herman

Stage two is where you add generation and run it in parallel. Solar on the roofs, or a CHP plant, or a gas turbine. The site now imports less, maybe exports sometimes. And the requirements change immediately. You need inverters or generators that can ride through grid disturbances rather than tripping unnecessarily, and you need anti-islanding protection so you don't back-feed the utility line.



Corn

So even stage two is doing real work.



Herman

Real work, and the failure pattern is that a site puts solar up, doesn't think about it, and then discovers that during a utility voltage sag their inverters all trip, and they lose all their generation right when they'd most like to have it. Ride-through settings are a real operational decision.



Corn

Stage three.



Herman

Stage three is where you add storage and a microgrid controller, and now you can island. So you need the seamless transfer, you need generation that can black start, meaning start without an external grid reference, and you need something that can balance supply and demand inside the island in real time.



Corn

Black start being the awkward one.



Herman

It's awkward because most generators aren't designed to start a dead grid. A big generator needs a reference to synchronise to, and if the grid is gone there's no reference. So you either have a smaller black-start capable machine, or a battery that can form the voltage and frequency reference the larger machines sync to.



Corn

Which is one of the quiet functions of grid-forming batteries.



Herman

It's the function that's least visible and most important. A utility-scale battery that can form a grid is the difference between an island that starts up and an island that sits there dead while everyone stares at each other.



Corn

And stage four is long-term autonomy.



Herman

Stage four is the hardest, and it's where you separate the people who've thought about it from the people who haven't. Long-term autonomy means you can run indefinitely without the public grid. Which means your generation has to cover your worst-case seasonal period, not your average.



Corn

And for a solar-heavy site, worst case is winter.



Herman

Winter, or a string of cloudy days, or a wind drought, or a mechanical failure at the wrong moment. So you either overbuild generation and storage substantially, or you accept load shedding during extended periods of low renewable output, or you keep a fuel reserve as a backstop.



Corn

Almost everyone keeps the fuel backstop.



Herman

Almost everyone, because the alternative is a capex number that nobody wants to look at. The honest answer is that long-term autonomy is a design decision with a cost, and the cost scales nonlinearly with how much of the year you want to cover with renewables only.



Corn

So the progression is more expensive and more operationally demanding at every step.



Herman

And the last step is where the curve goes vertical. Stages one and two are cheap additions to a normal building. Stage three is a serious engineering project. Stage four is a utility.



Corn

Now, this is the point where Daniel's second question gets interesting. Because for a residential community, that progression runs straight into questions that aren't technical at all.



Herman

Governance and money.



Corn

Let's start technical, then move.



Herman

Technically, a residential community microgrid needs a shared generation source. Rooftop solar on each house, or a shared solar array on common land, or both. It needs storage, at least one community battery, usually several. It needs a private distribution network, which is the awkward one, because most suburban neighbourhoods already have privately owned wiring from the utility's transformer to each house, and now you're proposing a second network behind it.



Corn

Easements, trenching, new ducts.



Herman

All of it. And it needs a control system that can decide at any moment whether the community is importing, exporting, islanding, or islanded. That's the microgrid controller again, just in a residential setting.



Corn

And in practice, the public grid stays connected.



Herman

Almost always, at least in the initial phases. The public grid becomes a backup or a supplemental source. Full disconnection is the rare case, and for most residential communities, is not what the business case actually justifies.



Corn

Which is one of the misconceptions worth naming. A community microgrid does not mean cutting the cable.



Herman

It doesn't. It means a community that has the option to island, and can island for some period, and can manage its own imports and exports more actively. That's a different thing from independence.



Corn

Ownership. What are the models?



Herman

There are a few that show up in practice. A homeowners' association is the simplest, because it already exists and already has governance. The problem is that an HOA board is made of volunteers, and a microgrid is a technical asset with maintenance requirements and regulatory obligations. You're asking a group of neighbours to run a small utility for an hour a month after their day jobs.



Corn

A cooperative.



Herman

A cooperative is more democratic, one member one vote, and it's the model a lot of rural electric co-ops already use in the United States. They've been doing exactly this, distributing power as a member-owned entity, for the better part of a century. So there's precedent, and there's a legal template.



Corn

The tradeoff is that cooperatives move slowly.



Herman

They move deliberately. Which is fine until you need to make a decision about replacing a transformer during a heat wave, and the next board meeting is in three weeks.



Corn

Special purpose entity.



Herman

An SPE is cleaner on paper. You set up a legal entity that owns the microgrid assets, the residents hold interests in it, and it contracts with a professional operator. That's how a lot of larger shared infrastructure gets structured. It's less democratic and more functional.



Corn

The fourth is a partnership with a private operator.



Herman

Which is where a company owns the assets, the community signs a long-term service agreement, and the operator carries the technical burden. The community gets less control and more competence. That tradeoff is real and the right answer depends on the community.



Corn

Now metering and payment, which is where things get legally weird.



Herman

They do. There are three main approaches. Master metering is the simplest, where the community takes one large service from the utility and resells electricity to its residents. Submetering is where each residence has its own meter, owned by the community, and the community aggregates the settlement with the utility.



Corn

And virtual net metering.



Herman

Virtual net metering is a regulatory construct that lets the output of a shared generation asset, like a community solar array, be credited to individual customer accounts on the utility's billing system. The utility does the accounting, the community just gets the credits distributed.



Corn

Which sounds like the cleanest one, and California and a handful of other states have versions of it.



Herman

They do, and it's very useful, but it depends entirely on the regulator allowing it. Change the net metering rules and the whole financial case shifts.



Corn

Which brings us to the regulatory problem, which is the biggest non-technical obstacle.



Herman

In most jurisdictions, reselling electricity is a regulated activity. If you sell a kilowatt hour to your neighbour, you may be acting as a utility and need a utility's licence. That's not an accident. It's because electricity regulation has historically been about protecting consumers from being bilked, and about ensuring the grid remains a public asset rather than a private toll road.



Corn

The community needs either an exemption or a partnership.



Herman

Or it structures it cleverly. If the community is a landlord and the electricity is bundled into rent, that's regulated differently than if it's selling power. If the community is buying a service rather than electricity, sometimes that works too. But anyone trying this should



Herman (continued)

have a lawyer on the team from day one.



Corn

And the economics.



Herman

The economics are really about what justifies the capital cost. A microgrid is a big up-front investment, and it's rarely cheaper than just buying from the utility on a pure dollar-per-kilowatt-hour basis. That's the first thing people get wrong.



Corn

So what justifies it?



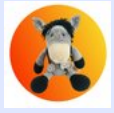
Herman

Resilience, mostly. Keeping the lights on during an outage has a value, and the value is largest for people who have something specific to protect. Medical equipment. Home-based work. A business that can't close. Communities in fire zones or flood zones or hurricane zones. And there's a soft value that's real, which is the sense of control, that the community is not just a passive customer of something happening fifty miles away.



Corn

The honest case is that a community microgrid is usually not a cost-saving move.



Herman

Not usually. Occasionally. If the community is in a place where the utility grid is unreliable, expensive, or overloaded, then the microgrid can pencil. There's an arXiv paper on community microgrid economics that makes exactly this point. The framework is about comparing the cost of the microgrid against the delivered cost of grid power plus the value of the resilience you gain. If either of those variables is bad enough, the microgrid wins.



Corn

If both are fine, it doesn't.



Herman

Then it doesn't, and you build it anyway because you want the resilience or the control. That's allowed. But you should be honest about which reason is doing the work.



Corn

Now the second-order question. If communities start doing this at scale, what happens to the public grid?



Herman

This is the one I think people should be worrying about more. If the wealthiest, most capable communities can afford to self-serve, they generate less demand, they pay less into the rate base, and the fixed costs of the public grid get spread over a smaller and poorer customer set. Which means rates go up for the people remaining, which makes the microgrid case look better for the next community, and so on.



Corn

A death spiral.



Herman

A slow one, with a lot of friction preventing it from fully playing out, but yes. It's the same dynamic that people worry about with rooftop solar in general. And the industry has been aware of it for a while.



Corn

The counter-argument is that private generation also reduces strain on the public grid.



Herman

It does, and if the private grid is willing to be islanded at the utility's request, the utility gets a benefit. But that requires the utility to have some control over the private asset, which means some kind of coordination agreement, which is a whole other layer of complexity.



Corn

The equity question is open.



Herman

It is. Which is one reason I think the governance questions matter more than the technical ones.



Corn

Speaking of governance.