Episode Book

EPISODE #5553

When 52 Volts Meets a 12-Volt Modem

PUBLISHED	25 September 2026 at 16:51 UTC
DURATION	26:51
LISTEN	myweirdprompts.com/episode/spare-parts-planning-home-network/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.2
BOOK GENERATED	19 October 2026 at 17:37 UTC

SYNOPSIS

A wrong AC adapter sent 52 volts into a 12-volt fiber modem, and the sizzle lasted a second and a half. That accident opens a bigger question: how do you decide which spares to stock when you can't afford to hold one of everything? This episode walks through the three-tier stocking model that managed service providers actually use, what the homelab resilience literature recommends, and how PJM — the largest wholesale electricity market in North America — thinks about spare equipment for the bulk electric grid. Along the way: why the cheapest component in your network is usually the one that dies, why a used-but-unreplaced spare is worse than no spare at all, and why "single node, move the drive" is a legitimate strategy that optimizes for something completely different.

TRANSCRIPT



Corn

Here's the thing about a fried modem. Everyone who hears the story says the same sentence. "You should have checked the adapter."



Herman

Which is true, and also completely useless.



Corn

Right, because the person who says it has never once been crouched in front of a cabinet at eleven at night with four barrels in their hand and one of them is the wrong one.



Herman

And they all fit.



Corn

They all fit. So Daniel sent us a whole thing about this. He was putting his network back together after Hannah had been painting a bookshelf and the gear got moved off it. He grabbed an AC-DC adapter, plugged it into the fiber modem, and fifty-two volts went into a device that wanted twelve. There was a slight sizzle. He actually said the sound made him wonder where it was coming from, which is the most honest part of the whole story.



Herman

By the time he'd placed the sound, the modem was gone.



Corn

Gone. And the adapter was the one for the Ethernet switch. So he's sitting there with no home internet until a replacement shows up, and he starts turning it over. He says he places a high premium on preparedness and failover, and this accident was trivially easy to make. Cramped cabinet, tired brain, split second, wrong barrel. And then he asks the



Corn (continued)

real question, which is the one we're actually doing today.



Herman

Which is where do you draw the line.



Corn

Where do you draw the line between preparedness and done enough. He could list every component in the chain that gets internet into his house. Modem, router appliance, server, access points, switches. He could buy a spare of every single one. He can't afford it, he doesn't have anywhere to put it, and most of it would rot in a box until it was e-waste. So how do actual planners, the people who think about risk and backup and disaster recovery for a living, decide what to hold and what to accept as risk?



Herman

And I want to say up front, he's not asking a naive question. He's asking the exact question that PJM asks about the bulk electric grid.



Corn

Let's start with the sizzle, though. What actually killed it.



Herman

Overvoltage, and not marginally. Fifty-two into twelve is four and a third times the rated voltage. Nothing in that device is built to see that. The failure happens in the power regulation stage, which is the part of the board whose entire job is to take twelve volts and turn it into the three point three and one point eight and whatever else the chips need. Feed it fifty-two and those regulators break down, and once they break down they pass the high voltage straight through to everything downstream.



Corn

Which is why there's no visible damage.



Herman

Nothing external, no. The barrel jack is fine, the case is fine, the LEDs are dark. It's all internal. The community reports of this are remarkably consistent. A pop or a sizzle, sometimes a smell of burnt circuitry, and a dead device. There's a case from years back of someone putting a twenty volt netbook supply into a twelve volt Huawei router. Same ending. Someone else put twelve volts into a nine volt modem and got about ten seconds before it made a noise and quit.



Corn

Ten seconds. So there's a window.



Herman

There's a window, and it's short enough that you will not notice it while you're doing something else. That's the thing. This isn't a slow degradation you can catch. It's a second and a half.



Corn

And the reason it's so easy is that the connectors are identical by design. A barrel jack is a barrel jack. The adapter fit the receptacle. The mismatch was printed in very small print on the side of a black brick.



Herman

The rules are simple and everybody half-remembers them. Voltage has to match exactly. Current, the amps, has to be at least the original rating, so a twelve volt two amp supply will happily run a twelve volt one amp device. And polarity matters even when the connector looks the same, because center-positive and center-negative use the same physical plug.



Corn

So the failure is a human factors failure wearing an electrical costume.



Herman

Entirely. And I'd push it further. The device was designed so that a wrong adapter can be inserted by a tired person in the dark. That's a design decision, and it's a bad one. There's no keying, no color coding required, no standard. Every other high-consequence connector in the world has a shape that stops you. This one doesn't.



Corn

So that's the mechanism. Which brings us to what you do about it, and that's where it gets interesting, because Daniel's instinct was to stock spares, and the professional answer to "which spares" is more structured than he'd expect.



Herman

It's a three-tier model. Managed service providers run it, and it's clean. Tier one is on the shelf. Tier two is same-day drop-ship. Tier three is never stock.



Corn

Define tier one for me, because the definition is where all the intelligence lives.



Herman

Tier one is parts that score high on two axes. High customer impact per day of downtime, and distributor lead time longer than twenty four hours. If a part fails and it hurts immediately, and you can't get one tomorrow, you hold one. Exactly one. Not two, not a drawer full.



Corn

And you replenish within forty eight hours of using it.



Herman

Within forty eight hours, because the moment you pull the spare is the moment you're unprotected, and the failure mode of a spare program is that you use the spare and never replace it and then you're worse off than when you started, because now you think you're covered.



Corn

That's the trap, isn't it. The spare you used and didn't replace is more dangerous than the spare you never had.



Herman

It's the worst of both. You spent the money and you have the false confidence.



Corn

So what's in tier two?



Herman

Anything a distributor can get to your site faster than you could drive a spare over. If CDW or Ingram or Anixter can have it there same day, holding it yourself is just capital sitting in a bin. You're paying for storage and depreciation to duplicate a service that already exists.



Corn

And tier three is the fun one.



Herman

Tier three is vendor-specific parts, low failure rate parts, fast-obsoleting parts, end-of-life models, anything the customer procured that isn't standard, and the whole category of just-in-case cables and adapters. The guide's phrasing on that last one is brutal. Every cable you don't actively use depreciates. The exotic adapter you buy just in case sits in a bin for four years and then you throw it out.



Corn

Which is exactly Daniel's e-waste worry, stated as a stocking rule.



Herman

And here's the part that made me laugh, and it's the whole episode in one line. The number one tier one recommendation in that guide is power adapters. Direct quote. Power supply failures are the number one silent killer for cheap network gear. Forty dollar part, thirty second swap.



Corn

So he didn't need a spare modem.



Herman

He needed a spare of the cheapest component in the chain and a label on it. The modem was never the vulnerable part. The adapter was.



Corn

That's a uncomfortable insight, because the thing that failed is the thing you fixate on. You buy a spare of the expensive box because that's the box that hurt.



Herman

And the expensive box is usually the most reliable thing in the rack. The cheap brick is the part that dies, and it's the part nobody thinks about.



Corn

Alright, so that's the small business answer. What does the homelab literature say, because Daniel's setup is closer to that. He's running a server, Home Assistant, his own fiber modem instead of the ISP box.



Herman

The homelab resilience guidance frames it in three layers. Data, hardware, service. And the key claim is that they only work together. Data redundancy without hardware redundancy means a single power supply failure still takes you offline. You've got perfect backups of everything and no machine to restore them onto.



Corn

Which is a real trap. People do the backup layer and feel finished.



Herman

The concrete tactics are what you'd expect. Dual network cards bonded together. Two managed switches running spanning tree. A backup internet uplink, which is the one most people skip. A UPS with a script that shuts the machine down gracefully instead of just dying. Dual power supplies on separate circuits, so one breaker trip doesn't take both. And spare drives, spare power supplies, spare cables.



Corn

And the warm spare.



Herman

The warm spare is a small preconfigured device sitting there ready to power on. Not a cold box in a closet, not a mirror of the whole system. Something that boots and takes over the one job that matters.



Corn

Which is a real design decision, because a warm spare is a maintenance commitment. You have to keep it patched.



Herman

You do. And the guide is honest about the other error, which is overengineering. Complex setups are fragile and difficult to manage. Start simple and add complexity only when justified.



Corn

Now here's where I want to push, because there's a counter-argument to all of this and it's a good one. There was a comment on a Hacker News thread about distributed systems that basically says the whole homelab high-availability project is a mistake.



Herman

I read that one. The person tried to build real high availability at home. GlusterFS for storage, corosync and pacemaker for clustering. And their conclusion was that it ends up being a lot of moving parts, expense, and headache.



Corn

And then the line I love. Guess what is fast? Single node. Move the storage drive to a similar spare small form factor PC. Two screws, takes about three minutes max.



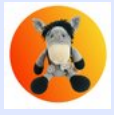
Herman

And their verdict on the whole enterprise. It's way too easy to overthink stuff.



Corn

Which is the exact opposite of the resilience guide, and both are defensible. One says layer everything, the other says keep one node and be able to move a drive in three minutes.



Herman

And the resolution isn't that one is right. It's that they're optimizing different things. The single-node person is optimizing for mean time to repair with a very simple procedure. The layered person is optimizing for staying up during the failure. Those cost different



Herman (continued)

amounts and they fail in different ways.



Corn

The single node approach has a hole, though, and it's the hole Daniel fell into. It works beautifully if you have the spare machine and the drive and you can do the three-minute swap. It does nothing if the failure is the thing you didn't think of.



Herman

Which is where the last piece of the homelab guide matters. Resilience isn't an accident, it's a design choice. And the design has to start from what you're actually trying to survive.



Corn

So let's take it up a level, because that's the question Daniel actually asked. Not what should a homelabber do. How do the biggest planners in the country draw this line.



Herman

PJM runs the largest wholesale electricity market in North America, and they have a document called the Spare Equipment Philosophy for Bulk Electric System Facilities. It's the clearest real-world answer I've found to Daniel's exact question.



Corn

And the first thing it does is sort equipment into three categories.



Herman

Major equipment, minor equipment, and common stock. And common stock, the stuff like control cable and fiber cabling and termination equipment, generally does not require sparing. They say it outright. No specific sparing methodology is necessary.



Corn

So the grid operator's answer to "should we hold spares for everything" is no.



Herman

Flat no. And the principle underneath it is the same one the MSP uses. Equipment critical to the integrity of the grid that's known to have long lead times should be supported by a spare. With particular focus on unique one-of-a-kind equipment.



Corn

One of a kind. That's the phrase that does the work.



Herman

Because a transformer that only exists in three places in the world is a twelve month problem if it fails. A control cable is a phone call.



Corn

So what are the actual lead times they're planning against?



Herman

Power transformers, thirty to sixty days. Transmission structures, about two weeks. Circuit breakers, one to two weeks. A capacitor bank back online within five days. Gas-insulated switchgear breaker mechanisms, one to two weeks. And they keep a minimum of three bushings in stock per transformer class, because a bushing is a small part that can strand a very large asset.



Corn

Thirty to sixty days for a transformer. That's the number that reframes the whole thing.



Herman

It reframes it entirely, because the spare isn't about convenience at that point. It's about whether the lights come back on this month. National Grid's strategic spares program puts the same logic in writing. Their strategic spares typically have lead times of six to twelve months, and they hold them partly to manage obsolescence.



Corn

Six to twelve months. So the spare is sitting in a warehouse for a year before it's needed, and that's still the right call.



Herman

And here's the sentence I think is the actual answer to Daniel's question. PJM writes that independent transmission operators maintain spare levels consistent with their risk tolerance for contingency events.



Corn

Say that again, because that's the whole episode.



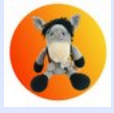
Herman

Spare levels consistent with their risk tolerance. Not with a formula. Not with an optimum. With their risk tolerance.



Corn

So the line isn't calculable. There's no equation that outputs the right number of spares.



Herman

There isn't. And I looked for one,. What you find instead is that the line is a policy decision that each organization makes for itself, and then backs into the inventory from there. Microsoft's reliability guidance does the same thing with criticality tiers. Each tier gets different availability requirements, and the requirements drive the business continuity plan. IBM states the ordering principle directly. Resiliency requirements should reflect



Herman (continued)

service level agreements and their recovery objectives, which in turn drive the technology. It should not be the technology driving the resiliency requirements.



Corn

Which is the sentence that would have saved Daniel a modem, honestly. Start from what you're trying to survive, then buy.



Herman

Start from the requirement. Most people start from the catalog.



Corn

So PJM uses probabilistic models to set quantities. Monte Carlo simulation, which they note you can do in Excel.



Herman

You can, and that's not a joke. They're modeling failure probability against lead time and deciding how many of a thing to hold. And they refresh the planning studies at least every five years, because the fleet changes and the failure rates change and last decade's answer isn't this decade's answer.



Corn

Five years. And the MSP's version of that is a twelve month no-use rule. Anything not used in twelve months is a candidate to de-stock.



Herman

Same logic at a different scale. The grid refreshes the study every five years, the small shop walks the shelf every quarter and de-stocks anything that hasn't moved in a year. Both are asking whether the thing we're holding still maps to a risk that still exists.



Corn

Now, the e-waste worry. Daniel raised it explicitly. He buys a spare modem, he upgrades to the next technology, the spare is trash. What does PJM do with obsolete gear?



Herman

They have a boneyard. Decommissioned equipment kept in storage, and they hold it up until the last of that style of equipment has been removed from their system.



Corn

So the rule isn't "keep it forever." It's "keep it as long as the fleet needs it."



Herman

Exactly that. The moment the last one of those units leaves service, the boneyard copy becomes scrap. Until then it's the only source of a part that nobody makes anymore.



Corn

Which validates Daniel's worry and also complicates it. He's right that the spare becomes e-waste. He's wrong that this makes it a bad purchase. It becomes e-waste on a schedule, and the schedule is set by when the in-service gear goes away.



Herman

And there's a real phenomenon here that PJM notes in passing. Solid dielectric cables can develop a permanent set after long periods of storage and become unusable. So a spare isn't inert. It can degrade on the shelf.



Corn

Which is a underrated point. The spare you bought three years ago and never tested might not work when you reach for it.



Herman

Which is why the MSP rule is a quarterly physical walk, not just a spreadsheet. You look at the thing.



Corn

Alright. Now the trade-off that I think is the honest heart of this. PJM says something that maps directly onto Daniel's situation, and it's about exposure.



Herman

They write that a transmission owner with very little underground transmission may not be able to justify holding those spare stocks. And then the next sentence. Unfortunately, this strategy can leave a transmission owner with minimal underground particularly exposed to a cable failure.



Corn

So the document names the exact trade. You can't justify the spare, and the thing you can't justify the spare for is the thing that will hurt you.



Herman

It's the same shape as Daniel's modem. He couldn't justify a spare modem for a device that had never failed. And the device that had never failed is the one that failed.



Corn

And I want to sit on that for a second, because it's not a gotcha. It's the actual structure of the problem. If you could predict which part fails, you wouldn't need a spare program. The whole reason you need one is that you can't.



Herman

The spare program is a bet against your own uncertainty.



Corn

There's a third path in that PJM document that I don't think gets enough attention, and it's the one Daniel didn't consider.



Herman

Shared spares.



Corn

PJM participates in sharing programs. There's one called STEP, there's RESTORE, and there's a national spare equipment list for transformers maintained at the Department of Energy level. Utilities pool the parts that are too expensive and too slow to duplicate.



Herman

Which is the answer to the affordability problem. If a transformer costs millions and takes a year, no single utility wants to hold three. So they hold them collectively and ship them to whoever needs one.



Corn

And the home equivalent is real. If Daniel has a friend running similar gear, the two of them holding one spare modem between them covers both houses for half the cost.



Herman

There's a local maker group version of that too. A shared shelf.



Corn

The catch is coordination. A pooled spare only works if the pool knows who has it and can get it there.



Herman

And if the pool is two people, it works fine. If the pool is twenty, you need a spreadsheet and someone who maintains it, and now you've built a small organization.



Corn

One more piece of vocabulary, because Daniel asked how planners think about this and this is the language they use. Recovery time objective and recovery point objective.



Herman

RTO is how long you can be down. RPO is how much data you can lose. And there are two more that matter for this conversation. Maximum tolerable period of disruption, which is the point past which the damage becomes unacceptable, and maximum tolerable data loss.



Corn

And the reason these matter is that they're set per critical function, not per system.



Herman

Per function. So for Daniel, the question isn't "how do I protect the network." It's "what's my RTO for home internet." If the answer is two days, he waits for the replacement and buys nothing. If the answer is four hours, he needs a spare modem and a spare router and probably a backup uplink, and he should know that going in.



Corn

Most people have never actually asked themselves the question. They just feel bad when the internet goes down.



Herman

They feel bad and they buy something, and what they buy is whatever failed last time. Which is how you end up with a spare of the expensive thing and no spare of the cheap thing that actually dies.



Corn

There's a storage point in the PJM document too, and it's the one that made me stop.



Herman

Spares shouldn't be stored where a single event can destroy both the in-service gear and the spare. They name collateral damage, flood events, and attack.



Corn

The spare goes in a different room. Ideally a different building.



Herman

The home version of that is real. If the spare modem is on the same shelf as the working modem, a fire takes both. If it's in the garage, a flood takes both. The spare has to be somewhere the failure can't reach.



Corn

Which sounds paranoid until you remember the whole point of the spare is the day something goes wrong.



Herman

The spare is insurance against a bad day, and insurance that's stored inside the thing it's insuring isn't insurance.



Corn

Alright. So Daniel's intuition was right, and the professionals agree with him, and they agree in a structured way. Stock the high impact long lead items. Don't stock the rest. Set the levels from your risk tolerance, not from a formula. And share what you can't afford to duplicate.



Hilbert

Forty dollars for the power supply. That's what they cost at the airport, roughly, for the voltage regulators on the runway lighting. I did the spare parts inventory at a small regional field for four years, and I want to correct one thing you said, because I think it's the thing that matters most and you both walked past it.



Corn

Go ahead.



Hilbert

You can plan your way to a spare for everything critical and still lose. We had a spare for every component on that lighting system. Every regulator, every photocell, every contactor. And it wasn't enough, because the failure we didn't plan for was a lightning strike that took out the entire control panel. The panel. The thing that all the spares plugged into.



Herman

The spares were fine and the chassis was gone.



Hilbert

The spares were fine and there was nothing to plug them into. And that's the part I'd want your listener to hear. Redundancy buys you time. It doesn't buy you invulnerability. Those are different products and people buy the second one expecting the first.



Corn

Say more about that, because I think that's a distinction worth naming properly.



Hilbert

When the strike happened, we had lights out on the main runway for about six hours. Six hours, in weather, at night. What got us through it wasn't a spare. It was that we had portable lighting units in the shed and a procedure for rolling them out, and we had a guy named Ray who knew which breakers to isolate so we didn't lose the taxiway too. The



Hilbert (continued)

spare parts never came into it. The plan did.



Corn

Six hours.



Hilbert

We had a spare for everything and it bought us nothing that night, because the failure was one level up from where we'd planned. The planning was still right. It just wasn't sufficient.



Herman

That's the lightning strike problem. You can't enumerate the failure, so you plan for the response instead.



Hilbert

You plan for the response. And the other thing, since you were talking about the boneyard. We had one. A shed behind the maintenance hangar, full of old equipment. Stuff that had been pulled off the system years before. And one night a regulator failed and there was no new one in the bin and no way to get one until Monday. Ray went out to the shed and came back with a regulator that had been sitting there since before I started. Thirty years old, at least. He put it in and it ran.



Corn

It ran.



Hilbert

Ran for another two years until they replaced the whole system. So your e-waste worry, the one about the spare becoming trash. That regulator was trash. It was trash for a decade. And then it was the only one on the property.



Herman

The boneyard isn't sentiment. It's a bet that the in-service fleet outlives the supply chain.



Hilbert

That's what it is. You keep the old thing as long as the old thing is still running somewhere. The day the last one comes off the system, you scrap it.



Corn

And the labeled bins.



Hilbert

The labeled bins are the whole job. I got yelled at once, properly yelled at, because a runway light went out during a storm and I couldn't find the spare voltage regulator because it wasn't in its bin. It was on a shelf two rows over because someone had borrowed the bin. The part was on the property. It may as well have been in another state. That's the lesson. A spare you can't find in ninety seconds is not a spare.



Herman

Which is the same as the quarterly walk.



Hilbert

It's the same thing. Anyway. I have a delivery coming that needs a signature and the window closes in about ten minutes, so.



Corn

Go.



Hilbert

I'll be back.



Herman

The thing I keep turning over from that is the lightning strike. Because Daniel's version of it is real. He could buy a spare modem, a spare router, a spare access point, and a spare server, and the thing that takes him down could be the power supply in the cabinet, or the fiber line itself, or a firmware update that bricks three devices at once.



Corn

Or the ISP.



Herman

Or the ISP, which no amount of home hardware fixes. And that's the correction Hilbert made. The spares are a buffer, not a wall.



Corn

Here's the misconception I want to put on the table, because I think it's the most common one and it's the one that costs people the most money.



Herman

Go on.



Corn

The belief that the goal is a spare for every component. That if you just buy enough, you're covered. And PJM says no in writing. Common stock items generally do not require sparing. There's no specific methodology needed. The people running the largest grid in North America looked at a category of parts and said we don't hold those.



Herman

The correction is that the line is a policy decision, not a formula. You set your risk tolerance, you set your recovery time objective, and the inventory falls out of that. It's not a number you can look up. It's a number you decide.



Corn

One more thing worth saying, and then I want to land this. The gap between home and enterprise practice is closing.



Herman

It is. The tools are the same tools. Tiered spares, recovery objectives, shared pools. Ten years ago a homelabber had never heard of an RTO. Now they're setting one for their Home Assistant instance because their spouse notices when the lights don't come on.



Corn

The frameworks scale down honestly. The MSP three tier model works for one person with one shelf. The boneyard concept works for a box in the garage. The sharing program works for two friends.



Herman

Which means Daniel's question has an answer, and the answer is that he already knew it. He just needed to hear that the professionals do it the same way.



Corn

This has been My Weird Prompts. Our producer is Hilbert Flumingtop, who is currently signing for a delivery.



Herman

If you enjoyed this one, leave us a review wherever you're listening. It helps people find the show.



Corn

Everything lives at my weird prompts dot com. We'll be back soon.