WHITE PAPER

INSIGHTS THAT POWER INNOVATION

LITEON®

Load-Bearing Ethernet

A Planning Guide to 802.3bt Standards, Costs, and Security

For Network, Facilities, and Security Teams

Version 2.0 — August 2026

Reader's Guide & Disclosure

A quick overview of the paper's purpose, usage, methodology, and disclosures.

About This Paper

This paper consolidates and supersedes LITEON's four-part PoE Strategic Series (July 2026), which covered emerging applications, deployment costs, security, and standard selection as separate documents. It carries forward the material that proved most useful and retires the rest.

How To Use It

- Sizing a refresh or writing a procurement spec — Section 3
- Building a budget or a business case — Section 4
- Accountable for uptime — Section 5 is the one that changes with PoE
- Security or risk — Section 6
- The Appendix is the working checklist, organized by deployment phase

Interest Disclosure

This paper is published by LITEON, whose product lines include power supplies and LED illumination components. Two of its recommendations touch those categories: Section 5's position that redundant power be treated as a design requirement rather than an upsell, and Section 2's treatment of PoE lighting. Both are stated as the wider industry states them, but read them knowing where they come from. The paper also argues in Section 3 against defaulting to higher-tier switching — a critique aimed at a category LITEON does not compete in, which is an asymmetry worth naming rather than leaving for a reader to notice.

Sources & Methodology

Numbered references appear at the end. Technical claims about IEEE standards are anchored to the standards themselves (802.3af-2003, 802.3at-2009, 802.3bu-2016, 802.3bt-2018) or to the Ethernet Alliance's published overview of them. Figures identified as vendor-reported reflect manufacturer claims rather than independent measurement. The 20–25% buffer margin used throughout is LITEON's own design guidance, labelled as such rather than presented as an industry standard. A short list of further claims we could not source to a citable original appears alongside the references. Adoption judgments reflect the market as of mid-2026.

Table of Contents

Reader's Guide & Disclosure 2
Table of Contents 3
Executive Summary 4
1. From Convenience to Infrastructure 5
2. What Is Actually Driving the Demand 5
3. Choosing the Standard: The Two-Clock Problem 8
4. Costing the Deployment 13
5. Failure Blast Radius 17
6. Securing What You Built 19
7. Conclusion 23
Appendix: The Consolidated PoE Deployment Checklist 24
References 26

Executive Summary

Power over Ethernet spent two decades as a background technology — the thing that quietly powers a desk phone or an access point without anyone thinking much about it. That era is ending. IEEE 802.3bt-2018 raised power sourcing equipment (PSE) output to 90W and guaranteed power at the powered device (PD) to 71.3W by using all four pairs of the structured wiring plant, roughly a sixfold increase over the original 2003 standard.1,2 That crosses a threshold. PoE can now genuinely replace, not merely supplement, traditional AC circuits for a widening class of building systems: Wi-Fi 7 access points, fanless edge compute, lighting, large displays, point-of-sale terminals.

This paper is a planning method for that shift, organized around the deployment lifecycle — choose the standard, cost the deployment, plan the failure domains, secure what you built.

Four things in it are not in the vendor comparison tables:

- The two-clock problem (Section 3). Your cable plant and your switches live on different lifespans — cabling warranties reaching 25 years against 5-to-10-year switch refresh cycles. "Futureproofing" is therefore two decisions with opposite economics and collapsing them into one is how organizations end up simultaneously over-spent on switching and under-built on cabling.
- A transparent TCO model (Section 4). Plenty of sources assert that PoE is cheaper than running AC circuits; almost none show the math. This one does, per device location, with every assumption labeled — including the one assumption that runs against PoE.
- Failure blast radius (Section 5). What a single switch failure now takes down, once cameras, access control, and phones share one power domain — and a method for deciding how much consolidation you can accept.
- Power as a security control, and its limits (Section 6). On PoE infrastructure the switch is both the network and the power utility. That gives incident response a lever it has on no laptop — and a trap, because power delivery and port state are separable in ways most operators assume they are not.

One discipline repeats across all four: specify against real device requirements, not nameplate maximums, and add margin. This cuts in both directions, and the paper is careful to separate them. Under-provision against the devices you know are coming and

you get degraded-mode deployments that look fine on paper. Over-provision against devices that may never arrive and you pay a premium for capacity your next refresh would have bought cheaper. Section 3 is about telling those two apart.

1. From Convenience to Infrastructure

When PoE was standardized in 2003, it solved a narrow problem: powering VoIP phones without an electrician-run outlet at every desk. IEEE 802.3af-2003 provided up to 15.4W over two pairs of Cat5e, and wireless access points followed a few years later.3 For most of the two decades since, PoE was a convenience — a way to avoid a second cable.

802.3bt changed the category. Pushing PSE output to 90W and guaranteed device power to 71.3W does not just extend the existing use cases; it admits a new class of device that previously required a dedicated electrical circuit. A fanless mini-PC, a Wi-Fi 7 access point running three radios, an LED lighting zone, a slim all-in-one display — these are building systems, not network peripherals, and they are now things an Ethernet cable can power. The cable, in other words, has become load bearing.

That shift has a consequence most planning processes haven't absorbed: the switch closet has quietly become a power utility. It is budgeted, procured, and refreshed like network gear, but a growing share of what depends on it would previously have been the electrician's problem. The sections that follow are organized around the four questions that change once that is true — which standard to buy, what it actually costs, what fails when it fails, and how to secure a port that carries electricity as well as data.

2. What Is Actually Driving the Demand

Wireless is the most demanding consumer

Wireless infrastructure is forcing a rethink of what "adequate" PoE provisioning means, and it is the most likely place an organization discovers its power budget is short. 802.3bt Type 3, offering up to 60W per port, is positioned for high-performance access points including Wi-Fi 7; Type 4, at up to 90W, targets the highest-draw applications. That distinction shows up directly in documented device behavior. Juniper's AP47 requires approximately 29W and "IEEE 802.3bt power for full functionality"; on 802.3at "the AP operates with reduced functionality" — the three Wi-Fi radios drop to 2x2:2, or 4x4:4 with only two radios enabled.4 Zyxel's WBE660S maps antenna configuration to power source in the vendor's own support materials: 4T4R on 802.3bt, 2T4R on 802.3at, and Wi-Fi disabled entirely on

802.3af.5 This creates a specific planning trap: a network can be Wi-Fi 7-ready on paper and still deliver Wi-Fi 6-level performance in practice, purely because the switch closet was budgeted for the previous generation's draw. Requirements vary device to device, so a blanket assumption across a multi-vendor deployment is not safe.

Limitations. The variability itself is the risk: "check the datasheet for your specific device" does not scale to a multi-site refresh. Type 4-capable switching also carries a cost premium that compounds across large port counts. And an access point can fail to power up correctly on a 60W switch that is not 802.3bt-compliant, even when the wattage looks sufficient — compliance, not wattage, determines compatibility.

Edge compute arrives at Type 4

Roughly 71.3W landing at the device is sufficient for fanless mini-PCs, point-of-sale terminals, and slim all-in-one displays — not just cameras and phones. This matters most in retail and hospitality, where POS terminals and edge analytics can share the switch already serving cameras and access points; in industrial and warehouse settings, where vision-based inspection nodes can be placed and relocated without an electrician visit; and across distributed branch or kiosk estates.

Limitations. Fanless devices running sustained high wattage face real thermal tradeoffs — an enclosure pulling near 71W continuously in a poorly ventilated closet is a different reliability proposition than the same enclosure at idle. There is also a cross-testing gap: edge compute vendors and PoE switch vendors do not always jointly validate against each other's hardware, so "rated for Type 4" on a spec sheet is not the same as validated interoperability with your switch. A device spec'd against Type 4 and deployed onto Type 3 infrastructure will underperform or fail to boot.

Lighting: the application with the longest track record

LED fixtures powered and controlled over Ethernet have the longest deployment history of anything in this section — Panduit was publishing dedicated design guidance on PoE as the fixed network foundation layer for smart buildings, lighting included, as far back as 2019.6 Lighting is also the application that most clearly makes the paper's wider point, because it is the one where PoE is unambiguously replacing an electrical circuit rather than

supplementing a network drop.

> **Limitations.** We are not able to source a defensible figure for how far adoption has actually gone. Claims that PoE lighting is now "mainstream in new construction" circulate widely and we could not trace them to independent measurement, so we do not make that claim here. What can be said with confidence is narrower and more useful: the design guidance exists and is mature, many PoE lighting deployments are associated with higher-power Type 3 or Type 4 infrastructure, although actual requirements vary by fixture design and lighting architecture, and the budget sits in facilities while the infrastructure sits in IT — which makes joint ownership from the design phase the single most important planning step, and the one most often skipped. Jurisdictional treatment of higher power classes under electrical code also varies, and belongs in an early conversation with the authority having jurisdiction rather than a late one.

On the watch list, not the roadmap

Single Pair Ethernet. IEEE 802.3bu-2016 introduced single-pair Power over Data Lines for 100BASE-T1 and 1000BASE-T1, letting a sensor take data and power on one thin cable in spaces four-pair cabling never fit.7 It remains meaningfully earlier in its adoption curve: interoperability work is still in progress, multi-vendor testing is thinner than the four-pair ecosystem, and deployment data at scale is sparse outside automotive and select industrial contexts. Watch-and-pilot unless an industrial automation or dense-sensor initiative is already underway.

Off-grid and solar PoE. Solar-plus-battery PoE has moved toward capable industrial equipment: a panel generates variable output, maximum power point tracking optimizes the harvest, a DC booster converts to stable 48V PoE, and a battery bank covers low sunlight. Commercial units add per-port scheduling and can shed lower-priority ports below a battery threshold. Treat the specifics with care — battery performance in extreme cold is a known constraint that vendor marketing rarely details, remote maintenance is a recurring cost that upfront pricing omits, and an unattended cabinet at a remote gate is a more accessible physical target than equipment inside a secured building (Section 6). The performance figures circulating in this category come from manufacturers marketing their own products and should be read as claims about specific SKUs rather than facts about the category. For any remote asset, the question to answer before deployment is port-level power budget and shedding priority under a multi-day low-sunlight scenario.

3. Choosing the Standard: The Two-Clock Problem

Every switch vendor publishes a comparison table of 802.3af, at, and bt, and most of them are accurate. What they rarely provide is a decision method — and when the answer is uncertain, they tend to resolve to the same advice: buy the bigger switch. That advice is incomplete, and it is worth noticing who gives it.

The standards at a glance:

Standard	Type	PSE Output (Max)	Power at Device (Min Guaranteed)	Pairs	Typical Devices
802.3af (2003), "PoE"	Type 1	15.4W	~12.95W	2	IP phones, basic cameras, badge readers
802.3at (2009), "PoE+"	Type 2	30W	~25.5W	2	PTZ cameras, Wi-Fi 5/6 APs, video phones
802.3bt (2018), "PoE++"	Type 3	60W	~51W	2 or 4	Wi-Fi 6E/7 APs, thin clients, displays
802.3bt (2018), "PoE++"	Type 4	90W	~71.3W	4	PoE lighting, edge compute, large displays

Four details from that table matter for the decision, and two of them are places the market routinely gets it wrong.

The device column is the one to compare against — the gap between PSE output and delivered power is cable loss, and vendors quote both inconsistently. Each standard is also backward compatible: an 802.3bt switch powers 802.3af devices safely through standard detection and classification, so buying up never strands existing devices. The reverse — a bt-only device on an af switch — is what fails.

Understanding "90W" and "100W-Class" PoE Terminology. Industry marketing often refers to "100W-class PoE," while IEEE 802.3bt Class 8 specifies 90W continuous PSE output. In practice, both terms are commonly encountered in vendor literature. The distinction is largely one of terminology rather than deployment planning, since procurement decisions should be based on IEEE classifications and actual device power

requirements. For specification purposes, IEEE 802.3bt Type 4 / Class 8 defines a maximum continuous PSE output of 90W and a minimum guaranteed PD power of approximately 71.3W. Procurement documents should therefore reference IEEE classifications and verified power budgets rather than marketing descriptions alone.

Type 3 is not necessarily four-pair. A Type 3 PSE "can either be capable of only 2-pair power, in which case it can support up to Class 4 (25.5 W PDs), or it can be capable of supplying 4-pair power in which case it can support up to Class 6 PDs (51 W)."2 While less common in enterprise deployments, IEEE 802.3bt permits Type 3 implementations that support only 2-pair power delivery. For this reason, procurement specifications should verify 4-pair capability explicitly rather than relying solely on the Type designation. Pair capability belongs on the procurement specification next to the Type — this is the single most avoidable way to buy "bt" and not get it.

Finally, Type and Class are different things. Type tells you the standard generation; Class (0–8) is how switch and device negotiate the actual allocation, which is why a switch with a modest total budget can serve a mix of low- and high-draw devices without reserving maximum power on every port. 802.3bt added Autoclass, an optional feature in which the PSE measures actual draw once at power-up and trims its allocation to match — useful where both ends support it, which is worth confirming rather than assuming.

Matching standards to devices

A working heuristic covers most of the device-matching question: if a device has a screen, a heater, a motor, or more than one radio, assume it has outgrown 802.3af. A plain desk phone runs on af; the same phone with a video screen wants at. A fixed indoor camera runs on af; add infrared illuminators, a heater for outdoor duty, or pan-tilt-zoom motors and it needs at or better. Wi-Fi 6 access points sit at 802.3at as the common default; Wi-Fi 7 access points frequently benefit from 802.3bt when full feature sets and maximum performance are required, though requirements vary by vendor and platform architecture, per Section 2. The heuristic is a triage tool, not a specification — the datasheet, not the standard's ceiling, is the planning number.

Two categories to handle deliberately

Vendor-proprietary high-power variants — Cisco UPoE at 60W and UPoE+ at 90W, negotiated via Cisco's own discovery protocols — predate or parallel 802.3bt and generally coexist with it, but they bind you to vendor-matched equipment on both ends of the cable. Price that in as a procurement constraint. Passive PoE is categorically different: fixed voltage applied with no detection, classification, or negotiation, common in budget wireless and surveillance gear. Because it supplies voltage regardless of whether the connected device negotiates or even supports PoE, a mismatched passive injector can damage the device outright, where IEEE-compliant equipment performs detection first specifically to prevent that. Passive PoE is not a standard to choose — it is a compatibility hazard to exclude.

The two clocks

Here is the question every comparison guide waves at and few analyze: should you buy more standard than you need today? The universal vendor answer — lean toward the higher wattage when uncertain — is not wrong so much as incomplete, because it treats the deployment as one purchase. It is two purchases with different lifespans, and the future-proofing logic inverts between them.

The cable plant is the long clock. Structured-cabling manufacturers sell tiered system warranties reaching 25 years — Panduit's Certification Plus program offers "a 15, 20 or 25-year standards-based, performance warranty," for instance.8 That is a warranty, not a measured lifespan, but it shows the horizon the industry designs around. Cable is disruptive and labor-dominated to replace, and the material difference between adequate and excellent is modest: Section 4's benchmarks put the Cat6-to-Cat6A premium near $75 per drop. A cable pulled today will plausibly serve three generations of switches and two of end devices. The economics therefore favor over-provisioning cable in new construction or major retrofits — provided you will capture its service life. Pulling Cat6 into a new building in 2026 to save $75 a drop is a decision the 2036 version of your organization pays for many times over.

The switch is the short clock. Enterprise PoE switches are commonly planned around a 7-to-10-year life, with budget models closer to 5 to 7. That inverts the over-provisioning math. If your five-year device roadmap does not include Type 4 loads, buying Type 4 capability today means paying the 802.3bt premium for capacity whose need, if it materializes at all, will probably arrive around your next refresh — when you could buy it at

that generation's price with that generation's feature set. Over-provisioning switches is cheap insurance only when the need is likely to arrive mid-cycle, inside the current switch's service life.

Stated as a rule: over-provision the asset you cannot easily replace, and right-size the asset you are going to replace anyway.

The counterweight vendor content rarely provides for a deployment whose five-year roadmap is phones, badge readers, fixed cameras, and standard access points, 802.3at is not a compromise. It is the correct answer. That scenario is common; it simply has fewer advocates.

> **Limitations.** This framing assumes normal refresh economics, and it assumes the bt premium is large enough to matter. We were not able to source that premium — see Section 4 — so treat the right-sizing recommendation as conditional on your own quotes rather than proven by ours. The framing also weakens when refresh cycles stretch well past ten years, when a known mid-cycle device transition is already scheduled (a Wi-Fi 7 rollout in year two of a seven-year switch life may strengthen the case for 802.3bt, depending on the architecture of the selected access points and the functionality required over their operational life), or when the organization's cost of a second procurement cycle is unusually high relative to hardware cost.

The decision framework

The framework takes three inputs — device roadmap confidence, cable plant condition, and switch refresh timing — and deliberately runs the cable decision and the switch decision separately.

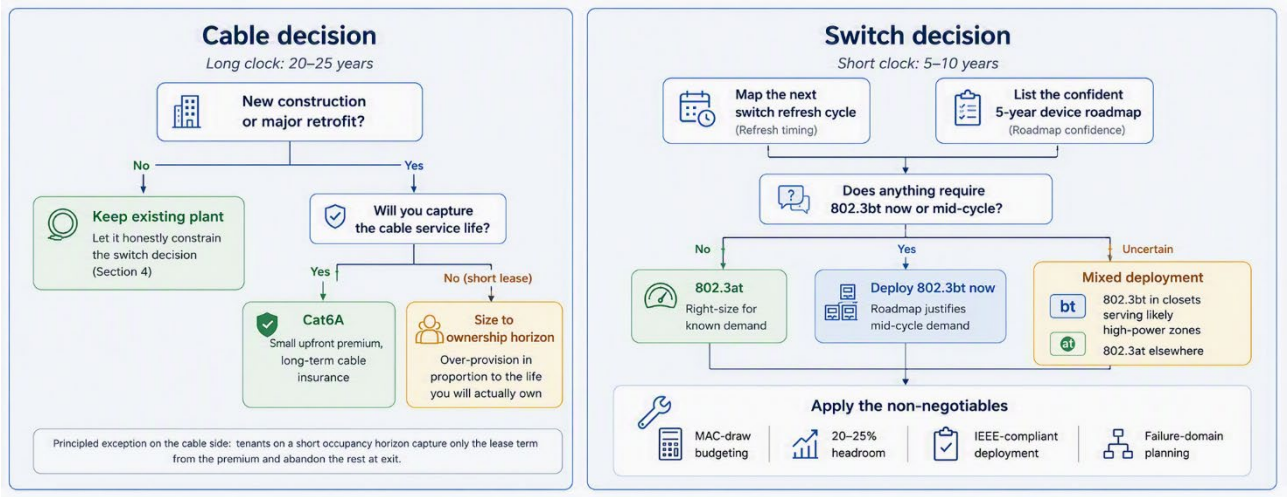


Figure 1. The two-track decision framework

— the cable decision and the switch decision run on separate clocks and are made separately.

One principled exception on the cable side: a tenant improving leased space on a short occupancy horizon does not capture 20 years of insurance from the premium — they capture the lease term and abandon the rest at exit. Over-provision cable in proportion to how much of its life you will actually own. And where genuine uncertainty survives planning on the switch side, a mixed deployment — bt in the closets serving likely high-power zones, at elsewhere — prices the uncertainty rather than buying it wall-to-wall.

Three worked scenarios

- Small office, 40 endpoints (phones, fixed cameras, Wi-Fi 6 APs; no expansion planned): retain the existing Cat6 plant, buy 802.3at switches
- Enterprise campus, Wi-Fi 7 refresh scheduled in year 2, cameras and access control expanding: Cat6A on all new pulls; bt Type 3 — four-pair capable, per the table note — in access-point-serving closets, at elsewhere
- New logistics facility, PoE lighting and edge compute in design scope: Cat6A throughout under the new construction rule; bt with Type 4 on designated closets, with failure-domain planning per Section 5 given that lighting and compute now share

switch power domains

> **Limitations** It treats one building on one clock. Multi-site estates will have closets at different points in their refresh cycles and should apply it per site, not once globally.

4. Costing the Deployment

Engineers who have deployed PoE at scale tend to describe the same failure pattern: everything works in testing, then months later devices start resetting intermittently, a switch closet runs hotter than expected, or a budget conversation reveals the cheaper rollout was not actually cheaper once cabling and redundancy are counted.

The power budget math nobody does upfront

Size PoE budgets to the maximum rated draw of each device, not its typical draw, because peak conditions — infrared illuminators at night, heaters in cold weather, motor activation — are exactly when underpowered devices fail. Treat per-port budget and total switch budget as two separate constraints that both have to clear.

LITEON's guidance is a 20–25% buffer margin: keep calculated peak load at or below 75–80% of the switch's total rated PoE budget. This is a design recommendation, not a requirement specified anywhere in 802.3bt, and we state it as ours rather than attributing it to an industry consensus we cannot document. It is where we have settled after weighing startup surge, seasonal peak, and headroom for growth against the cost of stranded capacity, and it sits within the range practitioners most commonly cite.

Choosing within the range. Use the 20% end where the device population is stable and well characterized, the environment is climate-controlled, and the five-year roadmap is firm. Move toward 25% where any of those are untrue — outdoor devices with heaters and illuminators, a roadmap with uncertain additions, mixed-vendor populations whose negotiation behavior you have not observed at scale, or closets where a later device transition is likely to land mid-cycle. If a deployment cannot clear 20% without a capacity upgrade, that is the signal to resize the switch rather than to shave the margin.

Worked example, illustrative. A 48-port switch with a 740W total PoE budget, populated with a mix of access points, fixed cameras, PTZ units, and door controllers totaling 628W of calculated peak draw, is running at roughly 85% of capacity. That is within nameplate —

112W remains — but it is past our guidance at both ends: a 20% margin caps the design target at 592W and a 25% margin at 555W. The deployment is over its own safety margin before a single device is added, and it has no room for the seasonal peak that would push the outdoor cameras toward their rated draw.

Oversubscription does not always fail cleanly. Depending on the switch, exceeding budget can mean a device is denied power outright, silently throttled, or the whole switch shedding load by port priority — and if port priorities were never configured, that shedding is effectively arbitrary. This is graceful degradation under oversubscription, distinct from the hard switch failure in Section 5.

Cable length, voltage drop, and derating

Voltage drop and thermal derating are two symptoms of one phenomenon — cable resistance and heat under sustained current — and are best treated together. The 100-metre channel limit remains the design boundary, but at higher PoE++ wattages, cable construction and bundle density affect how much of that channel is usable at full power. Cat6A conductors are typically 23 AWG against Cat5e's 24 AWG, and Cat6A's larger overall construction dissipates heat better in bundled runs — which is why it is the common recommendation for new PoE++ deployments and dense trays despite costing more per drop. Note that standard Cat6 is also commonly 23 AWG, so the advantage over Cat6 specifically comes from construction and shielding options rather than conductor gauge. Tightly bundled cables in poorly ventilated ceiling trays are the named failure pattern: heat cannot dissipate, resistance rises, and signal integrity and power delivery degrade together.

Interoperability traps

Legacy 802.3af/at devices generally negotiate safely on newer 802.3bt switches, because backward compatibility is built into the standard. The reverse direction is where problems appear: a high-power 802.3bt device connected to older or non-compliant hardware may fail to power up at all, because standards compliance — not rated wattage — determines whether detection and classification succeed. Note this is a different failure from the degraded Wi-Fi 7 radio modes in Section 2. There, negotiation succeeds and the device throttles to the power actually allocated; here, negotiation fails and nothing comes up. Both argue for verifying compliance between PSE and PD explicitly, but they present very differently in the field — one looks like underperformance, the other like dead hardware.

Infrastructure upgrade costs — the real TCO question

Plenty of sources assert that PoE is cheaper than running AC circuits; almost none show the math. Below is a transparent model, with every assumption labeled — including the one that runs against PoE.

Framing. This compares the incremental cost of reaching a single device location by each approach. It is not a whole-building comparison, and it excludes savings PoE offers beyond cabling — centralized backup power, remote reboot, fewer electrician callbacks — which would improve PoE's relative position further.

> **Limitations.** This model assumes that switching infrastructure, power budget, UPS capacity, and rack space are already available. These costs can materially affect the economics of a PoE deployment and should be evaluated separately

Benchmark inputs (US commercial, 2026)

Cost Item	Typical Range	Working Midpoint
Cat6 drop, standard commercial install	$150–$250	$200
Cat6A drop, standard commercial install	$200–$350	$275
Dedicated 120V circuit + standard outlet, commercial electrician	$250–$900	$575
Electrical permit (per circuit, where required)	$50–$350	$200

These ranges are compiled from general US commercial construction cost guidance and are directional. We are not citing a single authoritative source because none exists in comparable per-drop form. Treat the table as a model structure to populate with local quotes, not as a price list.

- **PoE path**: one Cat6A drop (~$275) carrying power and data to a switch that already has budget available.

- **Traditional AC path**: every device class in this paper is a networked device, so this path still needs a network connection — a data drop (~$200 for Cat6, since it carries no power) plus a dedicated circuit and outlet (~$575) plus, in many jurisdictions, a permit (~$200). Roughly $975 per location, before any panel capacity upgrade ($850–$4,000) if the panel is already full.

The exclusion that runs the other way. The PoE path above assumes a switch port with budget already available, which sets the per-port premium for PoE-capable switching to zero. That premium is real, and we could not source it (see below). Treat $275 as the floor of the PoE path rather than its full cost — PoE often reduces deployment cost, but actual savings depend heavily on switching infrastructure, existing cable plant, redundancy requirements, and the surrounding electrical environment. Organizations should validate project economics using site-specific quotations rather than relying solely on generalized models.

Where it holds, and where it doesn't:

- For a single isolated device in a building with spare panel capacity, the gap is roughly $275 against $975.
- For dense multi-device deployments, that per-device gap simply multiplies — 40 locations is roughly a $28,000 difference before any panel work. What genuinely compounds is risk on the AC side: each added circuit raises the odds of triggering the panel upgrade, a step cost PoE has no equivalent of, while PoE's shared switch cost amortizes across the group.
- The gap narrows or reverses when the existing cable plant needs a full Cat6A retrofit anyway, when switch capacity requires an unplanned upgrade specifically for the new load, or when a very small number of high-wattage devices never benefit from shared infrastructure economics.

On switch upgrade costs specifically. Enterprise 802.3bt switch pricing relative to 802.3at is not published in a clean, comparable per-port format. Vendor selection guidance consistently positions PoE+ tiers below 802.3bt tiers, but with no published percentage or per-port premium. We are deliberately not manufacturing a number. Get vendor quotes for your specific port count and power tier and treat that line as vendor-dependent rather than modelable from public benchmarks. Because Section 3's right-sizing recommendation

rests on this premium being material, that recommendation is conditional on your quotes confirming it is.

Regional labor rates vary meaningfully, so treat the midpoints above as a starting model, not a quote. The variables that most change the answer are device density per switch, existing panel and cable plant

5. Failure Blast Radius

Section 4 covered graceful degradation under oversubscription. This is a different failure mode: what happens when the switch itself fails outright, regardless of how well its power budget was managed.

Under traditional architecture, a camera's power comes from a local outlet or dedicated low-voltage supply, independent of its network switch. If the switch fails, the camera loses connectivity but may keep recording locally; if the power circuit fails, only devices on that circuit are affected. Under PoE consolidation, a single switch failure can simultaneously take down every camera, access-control reader, and phone connected to it, because power and data now share one failure domain.

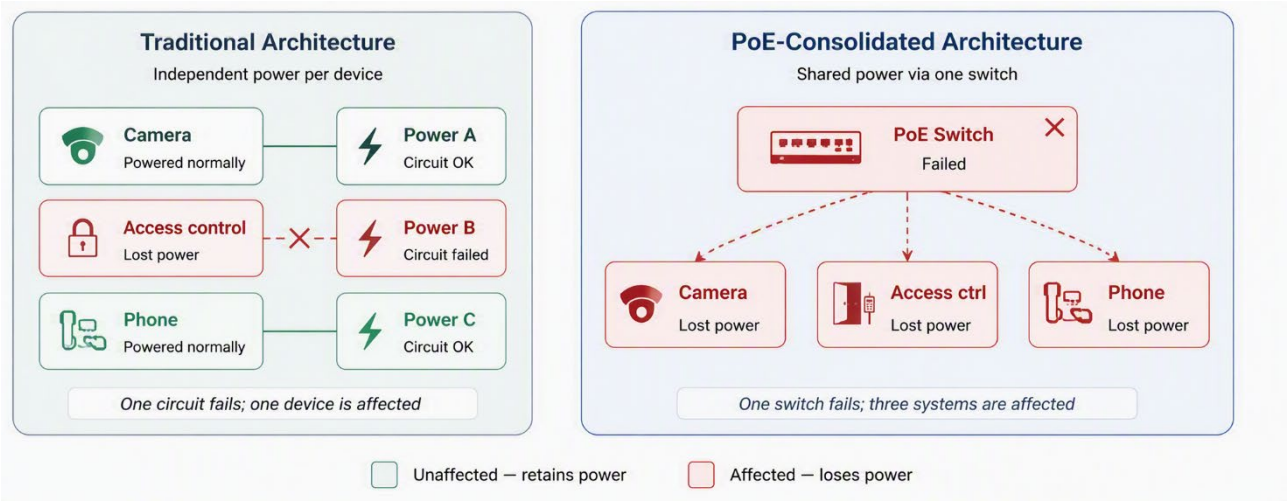


Figure 2. Failure Blast Radius

— traditional independently-powered systems against PoE-consolidated systems on a single switch

Stated explicitly: traditional architecture has more independent failure points, each smaller in scope, while PoE consolidation has fewer independent power systems but a larger blast

radius per failure. That is a trade of failure frequency for failure severity, not a strict downgrade. It is not a reason to avoid PoE — it is a reason to design the failure domain deliberately rather than discover it after an outage.

A method for sizing the domain

The tradeoff is easy to state and harder to resolve, so work it backwards from consequence rather than forwards from topology.

Four steps:

- Name the intolerable outage. Not "cameras are important" but a specific, falsifiable statement: "we cannot lose all badge readers on the north campus simultaneously," or "no more than one loading dock may lose surveillance at once." If nobody will commit to such a sentence, the failure domain has no design constraint and any topology is defensible — which is how zones end up on one switch by accident.
- Count backwards to a maximum devices-per-switch figure for each critical class. If no more than one of four docks may go dark, then no switch may carry cameras from more than one dock. That number, not port density or patch-panel convenience, drives the patching plan.
- Separate the two redundancy questions. Dual power supplies protect against PSU failure; a switch-level UPS protects against building power loss; distributing devices across chassis protects against everything else including software faults and line-card failure. They are not substitutes, and a dual-PSU switch on a single unprotected feed still fails completely when the feed does.
- Set the UPS runtime from the recovery objective, not the rack. The relevant number is how long the organization needs access control and life-safety-adjacent systems to survive an outage before staff can respond — which is a facilities and security answer, not a networking one.

In practice, this usually means not patching every important device into the first available ports on one switch — a pattern that creates thermal hot spots and means a single switch or line-card failure takes out an entire zone. For any switch carrying access control or life-safety-adjacent devices, redundant power should be a design requirement rather than an option.

6. Securing What You Built

Generic IoT security guidance treats PoE-powered devices as another category of connected endpoint and tells you to segment, authenticate, and monitor them. That guidance is correct as far as it goes and is well covered elsewhere, so it is summarized rather than repeated here. What it misses is specific to PoE: the power delivery itself is a security-relevant layer.

A scope note. This section describes vulnerability patterns and defensive architecture for a security-planning audience. It does not provide exploit techniques. Every vulnerability cited is public and documented in a vendor or government advisory; their remediation status varies, and one is included precisely because it has none.

Why these devices are different

Three traits set PoE-powered devices apart from the laptops and servers most enterprise security practice was built around. They are headless, with no user to notice something is wrong. They are frequently forgotten once installed, with security teams often lacking a complete inventory of what is connected. And their service life can run past the vendor's patch-support window — a mismatch examined at the end of this section. The trait that matters most, though, is physical accessibility: a camera on a building exterior or a sensor in a warehouse ceiling is reachable in a way a rack-mounted server behind a locked door is not.

The vulnerability pattern

CISA's advisory on an AVTECH IP camera (CVE-2024-7029) describes a flaw allowing unauthenticated attackers to inject and execute commands remotely, rated 8.8 on the CVSS v3.1 scale. It is also the example worth dwelling on: CISA records that "AVTECH SECURITY Corporation has not responded to requests to work with CISA to mitigate these vulnerabilities," and notes known public exploitation. There is no patch, and the recommended mitigation is network isolation.9 That is not an outlier to route around — it is this section's lifecycle problem arriving early, in a device that is still mounted and still powered.

Two further examples show the pattern classes. Bitdefender researchers identified a stack-based buffer overflow in the ONVIF request handler and a second overflow in the file-upload handler of Dahua cameras (CVE-2025-31700, CVE-2025-31701), either giving

unauthenticated remote code execution; they notified the vendor on 28 March 2025, patches shipped 7 July, and disclosure followed on 30 July.10 And a CISA advisory covering pan-tilt-zoom cameras from ValueHD, PTZOptics, multiCAM and SMTAV documents an authentication bypass at CVSS v3 9.1 and command injection at 7.2 (CVE-2024-8956, CVE-2024-8957), with no known public exploitation reported at time of publication.11

These are named because they are public and well documented, not because these manufacturers are unusually insecure. The same vulnerability classes appear across the IP camera and sensor industry broadly. Note also that beyond advisory-level findings, the more common and lower-effort risk remains default credentials and unmanaged shadow IoT — devices connected to the network that IT never inventoried, and that by definition are never evaluated against any of the above.

Segmentation and monitoring, briefly

The consensus does not need restating in depth: isolate device classes onto separate VLANs, apply 802.1X port-based authentication where devices support it, and treat MAC Authentication Bypass as a materially weaker control, since MAC addresses are trivially spoofed and it provides no cryptographic assurance of device identity. Segmentation is not a free upgrade — it adds VLANs to manage, rules to maintain, and new failure modes when a device legitimately needs to cross a boundary. Likewise, full behavioural-baseline anomaly detection at scale realistically requires a dedicated network access control (NAC) platform or a team with bandwidth to tune it. For a mid-size IT team without that, the achievable win is cruder but operable: maintain an accurate inventory and alert on unrecognized MAC addresses appearing on PoE-designated ports.

Power as a security control — and where it fails

Administratively shutting down a switch port is standard practice for isolating a compromised device, and on a PoE port it can do something no control does to a laptop: remove the device's electricity, not just its network access. On most platforms that is what happens — Cisco's own troubleshooting guidance states that "most Cisco switches turn off port power when the port shuts down."12

The word doing the work in that sentence is "most," and the exception is the more useful finding. PoE delivery runs on a state machine separate from the interface's administrative state, and several shipping

features deliberately decouple the two. Cisco's Perpetual PoE "provides uninterrupted power to connected powered device (PD) even when the power sourcing equipment (PSE) switch is reloading and booting up," while Fast PoE "remembers the last power drawn from a particular PSE port and switches on power the moment AC power is plugged in… without waiting for IOS to boot up."13 A port an incident responder shut

down is therefore re-energized on the next cold boot — before the configuration containing that shutdown has been applied. Other vendors ship equivalent always-on behaviour, in some cases enabled by default.

The correct reading is sharper than the simple one. On PoE infrastructure the switch is two utilities, and their control planes are not the same control plane. Power removal is a control to verify per platform, not a property inherited from shutting a port.

Two genuinely PoE-specific tools are commercially available today: remote power-cycling of individual ports, and per-port power scheduling. Scheduling has a narrow but real benefit — a device powered only during business hours has a smaller exposure window — though it does not apply to devices that must run continuously. What is not broadly available is power delivery gated directly on 802.1X authentication status, where a device receives no power unless it authenticates first. That appears in patent filings but is not something we found marketed as a standard switch feature; treat it as a direction the industry may move, not a control to plan around.

What to evaluate now. Confirm your NAC or 802.1X failure policy actually results in a full port shutdown for ports carrying high-risk device classes, not just a VLAN reassignment. This needs explicit confirmation because PoE power negotiation happens at the physical layer before authentication, so a port in the 802.1X unauthorized state typically still delivers power to whatever is plugged into it. Then confirm the shutdown itself de-energizes: check inline power state on an administratively down port, and disable perpetual, fast, or always-on PoE on any port where power removal is the security control — those features exist specifically to defeat this behaviour. Where the platform offers an explicit PoE-disable command, use it rather than relying on interface state.

Physical security

PoE devices are simultaneously a network endpoint and a physical power endpoint, routinely mounted in locations chosen for surveillance value rather than security — building exteriors, parking structures, warehouse ceilings, public lobbies. An exposed drop

or accessible enclosure is a physical path to whatever the device is connected to: an attacker with a few minutes of unsupervised access to a parking-lot camera can disconnect it and connect their own device to that port, inheriting whatever network access it grants.

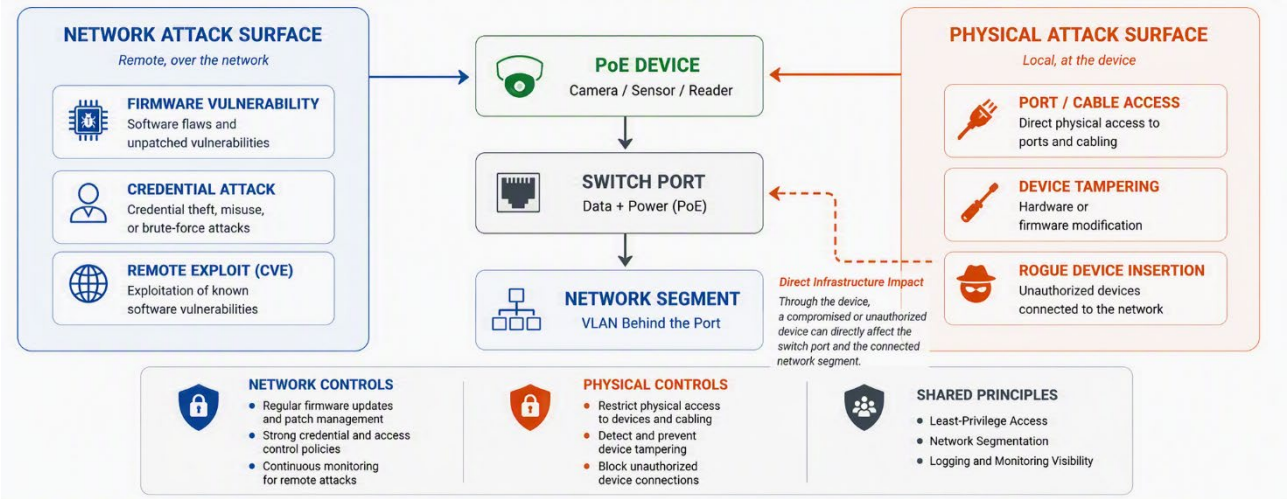


Figure 3. Physical Against Network Attack Surface

— a physical attacker at the device can bypass firmware and credential defenses entirely by inserting a rogue device into the same switch port.

This is where segmentation and port-level authentication earn their keep: an unauthenticated rogue device on that drop should get nothing — no VLAN access, and ideally no power. Treat as standard rather than optional: disable unused switch ports entirely rather than leaving them active and unpatched; apply port security to every port serving a physically accessible device; use tamper detection on enclosures where available; and physically secure any exposed patch panel, junction box, or injector cabinet.

A necessary caveat. These are deterrents, not guarantees. A determined attacker with time and basic tools can defeat most tamper detection and physical locks. Physical controls raise the cost and time an attack requires; they work as one layer alongside segmentation and port authentication, not as a standalone solution.

Firmware and lifecycle at scale

You cannot patch what you do not know you have, and shadow IoT is a persistent inventory gap in most organizations. The support mismatch is worth stating carefully.

Commercial-grade security cameras have a practical operational lifespan commonly cited in the 5-to-10-year range, and at least one enterprise-focused vendor publicly commits to five years of patch support after a product is discontinued — meaning a device bought near the end of its production run and then run for a typical lifespan can outlive its own vendor's support even under a relatively generous policy. That is an inference from one well-documented data point rather than an industry measurement; we do not have a broad survey, though vendors with shorter or less-documented commitments would plausibly produce a wider gap. The AVTECH advisory above is the harder case: not a support window that expired, but a vendor that did not engage at all.

The practical implication: firmware support lifespan belongs in procurement criteria rather than assumptions, and lifecycle planning should treat "still functioning" and "still receiving security patches" as two separate, independently tracked facts about every device in inventory.

7. Conclusion

PoE's story for its first two decades was convenience. Its next decade's story is that it becomes power and control infrastructure for categories — wireless densification, edge compute, lighting, remote assets — that had nothing to do with networking a decade ago.

The individual disciplines here are each documented somewhere on their own. What has been missing is treating them as one connected practice: choosing a standard is a two-clock economic decision, not a wattage lookup; costing a deployment is a real TCO exercise, not an assertion; reliability is a failure-domain design question the moment power and data consolidate; and security gains a lever — the power layer — that also has a failure mode most operators do not know about.

The compact version: separate the two clocks, over-provisioning the asset you cannot replace and right-sizing the one you will replace anyway. Run the cost numbers with your own inputs. Decide what outage you cannot tolerate before topology makes the decision for you. And verify, rather than assume, that the controls you are relying on do what they appear to do.

One last note on the advice this paper argues against. "When uncertain, buy bigger" is not always wrong sometimes the bigger switch is exactly right, and Sections 2 and 3 document the genuine forces pushing that way. But it is advice with a sponsor — and so, as the front matter notes, is some of the advice here. The discipline that serves a budget owner is not

defaulting up. It is knowing which clock each dollar is on, and asking who benefits from each recommendation, including ours.

Appendix: The Consolidated PoE Deployment Checklist

Organized by deployment phase and built to be used as a standalone reference.

Phase 1 — Select the standard

- List the confident five-year device roadmap; take maximum rated draws from datasheets, not the standard's ceiling
- Make the cable decision and the switch decision separately (Section 3's two clocks)
- Consider Cat6A for new PoE++ deployments and long-life infrastructure projects, particularly where bundle density or thermal performance is a concern
- Verify PSE/PD compliance explicitly; for Type 3 deployments, verify four-pair capability and supported power classes rather than relying solely on the Type designation. Some implementations may not provide the full power delivery commonly associated with four-pair Type 3 systems
- Treat any "100W per port" claim as a marketing description unless its IEEE classification and continuous power rating are explicitly stated
- Exclude passive injectors; price UPoE/UPoE+ as a both-ends procurement constraint
- Track Single Pair Ethernet interoperability work; pilot only if an industrial or dense-sensor initiative is underway

Phase 2 — Budget the power and the money

- Sum every device's maximum rated draw per switch, not typical or idle draw
- Confirm the result sits at or below 75–80% of the switch's rated PoE budget — LITEON's 20–25% buffer margin — choosing within that range per Section 4
- Check per-port class limits and total switch budget as two separate constraints
- Check bundle density and airflow in trays and conduit; confirm runs are within the 100m channel limit with margin
- Run the Section 4 TCO comparison with your own costs and density — including the data drop on the AC side
- Get vendor quotes for the 802.3bt premium — it is not modelable from public benchmarks, and Section 3's advice is conditional on it — and budget capacity

upgrades before rollout, not mid-project
- Confirm per-AP wattage against actual switch port budgets before any wireless refresh
- For PoE lighting, establish joint facilities/IT budget ownership and confirm code treatment with the AHJ, both from the design phase

Phase 3 — Plan the failure domain

- Write down the intolerable outage as a specific, falsifiable sentence
- Derive a maximum devices-per-switch figure for each critical class from it, and patch to that
- Distribute high-priority devices (access control, life-safety-adjacent) across multiple switches
- Treat dual PSU, switch-level UPS, and cross-chassis distribution as three protections, not substitutes
- Set UPS runtime from the organization's recovery objective, not from rack convenience
- Configure port priorities explicitly so shedding under oversubscription is not arbitrary, and define battery-shedding policy for any off-grid deployment at scale

Phase 4 — Secure the deployment

- Maintain a current inventory: make, model, firmware version, in-service date
- Alert on unrecognized MAC addresses appearing on PoE-designated ports
- Isolate PoE device classes onto separate VLANs; apply 802.1X where supported and treat MAB as a weaker fallback
- Confirm NAC/802.1X failure policy results in full port shutdown — data and power — for high-risk classes
- Verify empirically that an administratively down port is de-energized, and disable perpetual, fast, or always-on PoE where power removal is the control
- Confirm remote power-cycling is documented as an incident-response step, and evaluate per-port scheduling for devices that need not run continuously
- Disable all unused ports; apply port security to every port serving a physically accessible device
- Enable tamper detection on enclosures where available; physically secure exposed patch panels, junction boxes, and injector cabinets
- Track firmware support end-date separately from functional lifespan; set a replacement trigger on patch-support expiry, not visible failure
- Check procurement candidates against public advisory history — including whether

the vendor engages with disclosure at all

A note on this checklist's limits. Working through every item reduces risk without eliminating it. Site-specific conditions — electrical noise, firmware quirks in particular models, cable pathway constraints, vendor default configurations — can still produce failures a generic checklist cannot anticipate. Treat it as a floor for due diligence, not a substitute for a real risk assessment.

References

1. IEEE Std 802.3bt-2018, IEEE Standard for Ethernet — Amendment 2: Physical Layer and Management Parameters for DTE Power via MDI over 4-Pair. IEEE Standards Association.
2. Ethernet Alliance, "Overview of 802.3bt — Power over Ethernet Standard," 2018. ethernetalliance.org/wp-content/uploads/2018/04/WP_EA_Overview8023bt_FINAL.pdf
3. IEEE Std 802.3af-2003 and IEEE Std 802.3at-2009, DTE Power via MDI. IEEE Standards Association.
4. Juniper Networks, "Power Options for the AP47." juniper.net/documentation/us/en/hardware/ap47/topics/reference/ap47-power-requirements.html
5. Zyxel, "WBE660S WiFi 7 Access Point — Power Supply, special features and device behavior in case of power supply problems," Zyxel Support Campus. support.zyxel.eu/hc/en-us/articles/14843610154642
6. Panduit, "Power over Ethernet — Fixed Network Foundation Layer for 21st Century Smart Buildings," 2019. panduit.com/content/dam/panduit/en/solutions/solution-pdf/NI-ENT-PoEWP_FixedNetworkFoundationLayer.pdf
7. IEEE Std 802.3bu-2016, Physical Layer and Management Parameters for Power over Data Lines (PoDL) of Single Balanced Twisted-Pair Ethernet. IEEE Standards Association.
8. Panduit, "Certification Plus System Warranty." panduit.com/emea/en/legal-information/certification-plus-system-warranty.html
9. CISA, ICS Advisory ICSA-24-214-07, "AVTECH SECURITY Corporation AVM1203 IP Camera" (CVE-2024-7029), 1 August 2024. cisa.gov/news-events/ics-advisories/icsa-24-214-07
10. Bitdefender Labs, "Vulnerabilities Identified in Dahua Hero C1 Smart Cameras" (CVE-2025-31700, CVE-2025-31701), 30 July 2025. bitdefender.com/en-us/blog/labs/vulnerabilities-identified-in-dahua-hero-c1-smart-cameras

11. CISA, ICS Advisory ICSA-25-162-10, “PTZOptics and Other Pan-Tilt-Zoom Cameras” (CVE-2024-8956, CVE-2024-8957). cisa.gov/news-events/ics-advisories/icsa-25-162-10
12. Cisco, “Troubleshooting Power over Ethernet (PoE).” cisco.com/c/en/us/td/docs/switches/lan/catalyst3750 → Troubleshooting PoE
13. Cisco, “Configuring Perpetual PoE and Fast PoE,” Interface and Hardware Components Configuration Guide, Catalyst 9300 Switches. cisco.com → Catalyst 9300 Interface and Hardware Components Configuration Guide, IOS XE 17.1

Claims we could not source

Presented in the body as prevailing convention, not established fact. Each is load-bearing somewhere here, so a reader substituting their own judgement should start with these.

- The Section 4 cost ranges. Compiled from general US commercial construction cost guidance. No comparable per-drop source exists; it is a model structure, not a price list.
- Enterprise switch service life of 5–10 years. A planning convention consistent with common refresh cycles, not a measured statistic. It carries half the two-clock framework.
- The 802.3bt switch price premium. Direction is established; magnitude is not published in comparable form. Section 3’s right-sizing advice is explicitly conditional on it.
- Cat6A as the “consensus” choice for PoE++. Consistent across vendor and installer guidance, and with the thermal physics, but no standards body mandates it.
- PoE lighting adoption levels. Design guidance is mature and citable; adoption rates in new construction are not, and we make no claim about them.

Published by LITEON — Version 2.0, August 2026. Supersedes the four-part PoE Strategic Series, July 2026.