

RAPAX · WHITE PAPER

Fuse or Fold

The Future of Telecom Is Fusion

Run a bigger network with the team you already have.

Shawn Ennis — Founder & CEO of Rapax, a Citus Technologies company

September 2026

About the numbers in this paper. Operator-specific spend by team is not public — no carrier discloses what it pays for its NOC versus its SOC versus its contact center. So where this paper gives cost figures, they are one of two things, always labeled: **sourced** (attributed to a named source such as the U.S. Bureau of Labor Statistics, Gartner, or Omdia) or **modeled** (built from transparent staffing math with every assumption shown). We do not publish a single headline savings number, because a single number is a guess dressed up as a fact. We publish a method — and an interactive calculator so you can run it on your own numbers.

IN BRIEF

The argument in one page

Telecom has run on division of labor for forty years. Separate teams for separate jobs: a customer-experience team, a network operations center, a security operations center, an IT team — each staffed around the clock, each with its own ticketing, its own monitoring, its own idea of the truth. That model made sense when a router, a desktop, and an OC-48 BLSR ring each demanded a different human with a different skill set, and when money was cheap enough that you could throw headcount at every seat.

That world is gone. The 5G build-out did not move the revenue needle. Cheap money went with it. The people who know the old gear are retiring, and you cannot hire their replacements — nobody trains for a switch whose maker went out of business a decade ago. Meanwhile you are being asked to run more network, defend against more threats, and answer more customers, with less.

There is one thing that is finally strong enough to sit across all of those jobs at once: **AI**. Not to remove the humans — customers and executives still want a human on the other end, and your best engineers are more valuable than ever — but to absorb the repetitive volume that has been burning those humans out on the 3 a.m. shift. That is what makes consolidation possible for the first time.

This paper makes the case for the **Operations Fusion Center (OFC)**: one AI-augmented team that triages every signal — from any customer and any piece of infrastructure — resolves what it can, tickets and escalates the rest, and hands the genuinely hard problems to a daytime engineering function staffed by the best people you used to waste on overnight monitoring. It covers what the separate-team model actually costs, what fusing it saves, an honest range for the bottom line, and a realistic path to get there: roughly **90 days** for a small operator, **six months** for a Tier 2, and **twelve to eighteen months** for a Tier 1. It ends where the title starts. Fuse, or fold.

The past needs to be the past

I have spent twenty-five years in the guts of telecom operations. In the 2010s I helped do something that felt radical at the time: take fault management, performance management, topology and inventory, and service management — four disciplines that lived in four different tools, run by four different groups — and fuse them into a single unified service-assurance layer. One view. One source of truth. It worked, and the industry followed, because the alternative — swivel-chairing between systems that each held one-quarter of the picture — had become indefensible.

That was fusing the **tools**. This paper is about the next fusion, the one the tools made possible but couldn't finish on their own: fusing the **teams**. And the reason we can attempt it now, when we couldn't then, is that AI and AI-native knowledge management have finally become the common denominator across work that used to require four separate kinds of expert.

Look at how the old model was built. You hired a NOC technician to watch the network. You hired a SOC analyst to watch for attackers. You hired contact-center agents to answer customers. You hired an IT team to keep the internal lights on. Each was a distinct craft with a distinct culture, distinct tools, and — critically — a distinct budget owner. A router alarm, a desktop ticket, and a fiber ring fault were genuinely different problems requiring genuinely different people. Specialization wasn't a mistake. It was the correct answer to the question the era was asking.

The era stopped asking that question. Four things changed at once.

The money changed. The 5G build-out was sold as a growth engine and did not deliver one. Even as total service revenue grows, blended ARPU is going the wrong way: Omdia projected average revenue per user falling 4.2% between 2022 and 2027, and described the apparent uplift from unlimited 5G bundles as an illusion. Analysys Mason reached the same verdict — that 5G on its own is not enough to drive ARPU growth. Capital intensity peaked around 18% in 2022 and capex contracted in 2023 for the first time since 2017 (Dell'Oro). Global telco revenue slipped to a decade low in 2023 (GlobalData). And the debt that funded all of it got more expensive the moment rates rose. The growth era ended and the efficiency era began — not as a slogan, but as a balance-sheet fact.

The headcount already started falling — bluntly. AT&T's workforce went from roughly 280,000 in 2017 to about 141,000 by early 2025 — nearly 139,000 fewer people. Verizon's September 2024 SEC filing disclosed about 4,800 separations under a single voluntary program by the end of March 2025, against a severance charge of \$1.7–1.9 billion. Equipment makers cut too; Nokia announced up to 14,000 jobs in October 2023. The industry has already accepted that the answer to flat revenue is a smaller cost base. What it has mostly done so far is cut with a blunt instrument. Fusion is the version that cuts with a scalpel — and redeploys instead of only removing.

The knowledge is walking out the door. Telecom is hit harder than almost any industry by the retirement wave, because so much of the network is old and so much of the expertise is undocumented and concentrated in one or two veterans per site. More than half of the professionals who understand legacy systems have already retired, and fewer than one in five new hires can work on them (Techolution). You cannot hire a twenty-five-year-old who knows your SONET/SDH ring or your TDM switch. The vendor is

gone; you buy spares on eBay. When that person retires, the knowledge retires with them — unless you have captured it somewhere a machine can serve back on demand.

The cross-domain hire doesn't exist. A router, a desktop, and an OC-48 BLSR ring still require different skills. The difference is that you can no longer find, afford, or retain enough humans who span them. The labor math is unforgiving, and we'll get to it. The point here is simpler: the specialization that was once a strength has become a staffing trap.

Here is the turn. For the first time, there is a technology that can sit across all four of those jobs at once — reading a customer's message and a device's syslog with the same engine, triaging both, resolving the routine, and knowing when to call a human. AI is the common denominator the old model never had. It does not eliminate the need for skilled people. It changes **where** they belong — off the overnight monitoring shift and onto the engineering problems only a human can solve. And it lets you keep the thing customers and executives will never stop wanting: a person, on the other end, who actually knows what they're doing.

The past needs to be the past. Not out of nostalgia's opposite, but out of arithmetic.

PART 2

What fusion means: the Operations Fusion Center

The word to reach for is already in the market, and it's spreading fast — which is good news, because it means you are not asking your board to bet on something no one else believes.

The term is **fusion center**. It began in military and law-enforcement intelligence: a single hub where signals from many sources are merged so that an analyst calls **one place** instead of hunting for scattered fragments. Cybersecurity borrowed it as the **cyber fusion center** — an evolved SOC that unifies threat intelligence, detection, incident response, and sometimes IT operations and fraud into one coordinated unit. And in 2026 the network joined in. Splunk has predicted that AI will merge the NOC and the SOC into a “fusion center,” letting leaders redeploy talent instead of maintaining large siloed teams. Kyndryl launched a Cyber Defense Operations Center in February 2026 that fuses network and security operations into one operating model. Cisco and Splunk stood up a unified SOC and NOC at Mobile World Congress 2026. Gartner forecasts that by 2026, 75% of organizations with operational-technology environments will merge their NOC and SOC functions. Nemertes has argued the logic for roughly seventeen years.

So the direction of travel is settled. The industry is converging on the fusion center. **Our argument is that it does not go far enough — and that where it stops is exactly where telecom can least afford to.**

Every version above fuses the network with **security**. It is defense-first: fuse so you catch the threat faster. Look at what's missing from all of them. The **customer** is not in the room. The person who calls in, the ticket that opens, the service experience that determines whether they churn — none of it is a co-equal signal in a “fusion center” designed to stop attackers. General **IT** shows up only as a data feed, never as a fused function. The industry is fusing to defend. Telecom has to fuse to **survive** — which means the customer and the running of the network are not secondary signals to be bolted on later. They are the point.

That is the Operations Fusion Center. Not a security fusion center with a network feed attached. **One team where the customer, the network, the threat, and IT are equal citizens — anchored on running the business, not only on defending it.**

Concretely, the OFC works in three layers:

One front door for every signal. Customer messages and infrastructure signals — a support request, a syslog event, a security alert, a helpdesk ticket — arrive at a single AI-driven intake that triages across all four domains against one source of truth. No more “not my problem.” No more a fault that the NOC sees, the SOC needs, and neither hands to the other in time.

AI resolves the routine; humans keep the touch. The fusion layer handles Tier-1 and much of Tier-2 automatically: it deflects and resolves the repetitive volume, opens and enriches tickets, and escalates with full context when it can't. A smaller human team stays on the line for the moments that need a person — because customers and executives want a human, and that never changes. What changes is that the human is no longer buried under 3 a.m. noise.

Engineering moves to daylight. The genuinely hard problems — the ones that need judgment, design, and deep expertise — go to a daytime (8x5) engineering function reconciling during business hours. This is where your scarcest, most senior people belong. Not watching a screen overnight. Solving what only they can solve, with a life that doesn't drive them out the door.

Two things this is **not**, because both misreadings will come up in the room:

It is **not** the “lights-out NOC.” This paper is close to a rebuttal of that fantasy. Fusion keeps people — deliberately — and moves the best of them to higher-value work. The goal is a bigger network run by the team you already have, not an empty room.

And it is **not** merely the NOC-plus-SOC merger the market is already doing. It is one step beyond it, on purpose. The customer is in. IT is in. And — to be clear about the end state — so is security. Completely.

The one sacred cow this paper is here to challenge

The instinct that security must be an organizationally separate team is the strongest objection you will hear, and it deserves a straight answer rather than a dodge. Here is the straight answer: **almost all of it is orthodoxy, not law.**

We checked the frameworks that people invoke, one by one:

- **SOX** does not require a separate security team. It requires management to assess internal controls and an auditor to attest to them; separation of duties enters only as a control activity companies choose under COSO — an implementation choice, not a mandate.
- **PCI-DSS** requires separation between development/test and production environments, explicitly satisfiable with role-based access. Not a separate team.
- **NIS2** requires board approval and oversight of cyber-risk measures and makes leadership accountable — and it explicitly permits delegation of operational security. It does not require an independent function.
- **ISO 27001 and the NIST Cybersecurity Framework** require that conflicting duties be segregated, with compensating controls (logging, supervision, audit trails) available to small teams. Neither requires a standalone security org.

There is exactly **one** genuine constraint, and it is narrow. Under CALEA, a lawful intercept may be activated only with the affirmative intervention of a designated carrier officer or employee, and the carrier must name a senior person accountable for intercept security and file a system-security plan with the FCC. That is real, and it is non-negotiable. But notice what it actually requires: **controlled access and a named, accountable individual — not a separate department.** It is fully satisfiable inside a fusion center by ring-fencing intercept activation to CALEA-designated personnel with role-based access. Everything else — the belief that the SOC must report through a different chain to a different chief with a different budget — is culture, turf, and habit wearing the costume of compliance.

So let's be honest about what will actually happen, because pretending otherwise would insult the reader. **Politics will fight this.** The CISO and the CIO rarely report to the same person; the budgets are separate; the missions and the mindsets differ. In practice, most operators will delay full fusion, and many will water it down. That is fine. This paper is not a prediction of what the average telecom will do next quarter. It is a description of what **absolute** looks like — the pure form of the idea, so you can decide with open eyes how far toward it you intend to go, and so that the parts you leave separate are separate by **choice** and by genuine constraint, not by inertia you never examined.

PART 3

The solution: Rapax

An Operations Fusion Center is an operating model, not a product. But it needs a platform built for it — and most of what's sold to telecom is not. Legacy assurance stacks are tools you integrate around; observability platforms are repurposed IT monitoring pointed at a network. Rapax is different by design: **the agents are the architecture, not a layer bolted onto legacy systems.** That is the difference between a fusion center you can actually staff and a dashboard that still needs four teams behind it.

Rapax was built to do exactly the three things the OFC requires:

Collect from any infrastructure — fast. The fusion model only works if every signal actually reaches the front door, including the signal from thirty-year-old gear. Rapax's integration engine, **Bruce**, builds and deploys connectors — syslog, SNMP, webhooks, REST — to bring in data from anything, new or ancient. This is what makes AI-enabled migration realistic rather than a multi-year project: at one mid-market fiber operator with 600,000+ subscribers, Rapax stood up and displaced the incumbent monitoring in a proof-of-concept measured in **hours, not months.**

Interface with any customer and absorb the Tier-1 volume. **Frank** is Rapax's frontline agent: it validates intake, deflects and auto-closes the routine, and escalates with full context when it shouldn't act alone. And the proof point here is not a slide — it's production. Rapax runs Frank on its **own** support desk, and since May 2026 it has deflected roughly **40%** of tickets and auto-closed roughly **15%** — real numbers from a live queue, not a vendor benchmark from a tuned lab. (Independent industry results for AI triage are wide-ranging and often lower than vendor claims; we quote our own production because it's the number we can stand behind.)

Give the team memory it can trust. This is the piece that turns the retirement problem from a threat into a non-event. Rapax's AI-native knowledge engine, **Wade**, captures runbooks, procedures, and vendor documentation and serves them back on demand — with a confidence score attached to every answer, so

the team knows when to trust it and when to verify. The result: **your fusion team needs to be expert in Operations, not in a dead vendor's manual.** When the last engineer who knew the old switch retires, the knowledge doesn't leave with them.

Put together, that is a platform on which one team can actually run the customer, the network, the threat, and IT — which is the whole point.

PART 4

The economics

This is the part the CFO reads first, so it gets no hand-waving. We'll build the cost of the separate-team model from the ground up, show where fusion creates savings, and give an honest range for the bottom line — with a worked example and, at the end, the formulas that power the calculator on our site.

Why round-the-clock is so expensive

Start with the constant that governs everything: a seat covered 24/7/365 is not one person. A week has 168 hours; a full-time schedule is 40. That's **4.2 people** to cover a single seat before anyone takes a vacation — and realistically **5 or more** once you add a relief factor for the 15–20% of time real humans spend on PTO, sick leave, and training. Three people on duty at all times is not 3 hires. It's 12 to 15.

Now price each of those people fully loaded. Salary is only part of it: benefits alone run about 30% of total compensation (BLS), and once you add payroll taxes, workers' comp, equipment, software, facilities, and recruiting, the fully-loaded figure lands at roughly **1.25 to 1.4 times base salary** (a long-standing rule of thumb). Representative U.S. medians give you the base: information-security analysts around \$129,000 (BLS), network/systems administrators around \$99,000 (BLS), NOC technicians in the mid-\$50,000s (market data), and customer-service representatives around \$44,800 (BLS — which, tellingly, projects a 5% decline in these roles through 2035, attributed directly to automation).

Stack the multipliers and one number tells the whole story: a **single SOC seat covered 24/7 costs an estimated \$480,000 to \$730,000 a year** (industry modeling), and a minimum viable round-the-clock SOC needs 8 to 12 people. Staffing is roughly 65–70% of total SOC cost. Repeat that structure for the NOC, the contact center, and IT, and you begin to see the size of what four separate 24/7 towers actually costs.

The cost you don't see on the org chart: burnout and churn

The overnight shift doesn't just cost salaries — it costs **people**, repeatedly. In security operations, 71% of analysts report burnout and roughly two-thirds are considering leaving within a year (Tines); turnover runs 20–30%, and replacing an analyst costs 50–200% of salary. Contact centers are worse: 30–45% annual turnover, average tenure around 13–15 months, and \$10,000–20,000 to replace each agent — in a function where labor can be up to 95% of the cost. Every one of those departures takes context and knowledge with it. This is money bleeding out sideways, and it is precisely the human cost the fusion model relieves by taking the repetitive overnight volume off people's plates.

The lever hiding in your outsourcing line

Much of the Tier-1 work is outsourced today — contact-center BPO at roughly \$6–16 per agent-hour offshore, \$10–22 nearshore, and \$25–45 onshore; managed security at \$5,000–25,000 a month for a mid-market retainer. The uncomfortable truth about that spend: outsourcing's realized savings routinely disappoint — one large study across 368 buyers put the average at just 15% (ISG). Once AI can deflect the routine in-house, that money becomes a lever. You bring Tier-1 ops and IT back inside, automated, at lower cost and with better control — instead of paying a third party a thin margin to do it badly.

Where the money is, in context

None of this is a rounding error. Labor is about 17.4% of total telecom OPEX and roughly 22.5% once you exclude depreciation (GlobalData), and network operations has climbed to about 18.4% of OPEX. Globally, telco labor cost ran about \$263 billion across 4.5 million employees in 2023. The efficiency era is already reaching it — labor OPEX finally fell 2.2% in 2024 (Omdia) as the first workforce transformations landed. Fusion is how you make that a deliberate strategy rather than an annual round of blunt cuts.

The synergies of fusing

Combine the four towers and the savings come from five places, not one:

1. **Collapsed coverage.** One 24/7 front door instead of four means you stop paying the 5x-per-seat multiplier four times over. This is the single largest lever.
2. **AI deflection of the routine.** Automating Tier-1 across all domains removes volume rather than people wholesale — cross-domain triage can deflect 60–80% of routine tickets, and event correlation cuts alert noise 80–95% (industry data; a 1,000-device environment can go from tens of thousands of raw alerts a month to fewer than 500 actionable incidents).
3. **Insourced Tier-1.** The outsourcing line shrinks as automated in-house handling replaces BPO.
4. **Consolidated tooling.** Separate ticketing, monitoring, and inventory become one — the savings our companion paper, *The Rationalization Window*, quantifies in detail.
5. **Redeployed talent.** Your best people move from overnight monitoring to daytime engineering, where their output is worth far more — and where they're far more likely to stay.

The bottom line — as a range, with the split shown

We will not hand you a single percentage, because the honest answer depends on your seat counts, your blended cost, and — crucially — a choice only you can make: how much of the freed capacity you **redeploy to engineering** versus **remove from the budget**. So here is a worked, clearly-modeled example, and then the tool to run your own.

Modeled example — a Tier 2 operator. Suppose four 24/7 towers (CX, NOC, SOC, IT) with three covered seats each: 12 seats. At 5 FTE per seat, that's 60 people. At a blended fully-loaded cost of \$120,000, that's **~\$7.2 million a year in coverage labor alone** — before tooling and BPO.

Apply the fusion model — 60% deflection of Tier-1 volume, engineering shifted to 8x5, Tier-1 pulled back in-house — and a defensible estimate is that you take **30–45% off that run-rate: roughly \$2.2M–\$3.2M a year**, before platform cost, with tooling and BPO savings on top.

But that headline splits two ways, and you control the split. If your goal is a leaner budget, most of it drops to the bottom line. If your goal is capability — the people-first path this paper argues for — much of it is redeployed: the same money, the same people, now aimed at engineering and growth instead of overnight firefighting. Both are wins. They are different wins, and you should choose deliberately.

Your numbers will differ. That's the point of the calculator — put in your seats, your blended cost, your coverage model, and your target deflection rate, and see your own range with your own redeploy-versus-reduce split. The formulas are in the appendix; the tool is on our site.

PART 5

A realistic path

The vision is absolute; the rollout is not. You don't reorganize the company on day one — you prove the model, then move the org chart to match. Three stages, sequenced so that the risky part comes only after the safe part has earned it.

Stage 1 — Unify the pane of glass and Tier-1 triage, before you touch the org chart. Bring every signal to one front door, deploy AI-driven cross-domain triage, and stand up AI-native knowledge (Wade) so the team's memory is captured before anyone retires. Use Bruce to roll monitoring across your infrastructure via AI-enabled migration — including the legacy gear — so nothing is left blind. This is the lowest-political-cost, highest-return move, and it changes no one's job yet. Set clear gates before proceeding: **≥60% deflection of routine tickets and ≥80% reduction in alert noise**. If deflection stalls below ~40% in your environment, pause and tune the knowledge base before you go further — the automation layer isn't ready, and cutting headcount against it would be a mistake.

Stage 2 — Stand up the fusion team and move your best people to daylight. With the pane of glass unified and Tier-1 automated, form the fusion team for Tier-1/Tier-2 across all domains, and — this is the heart of it — **convert the top people from each of the old teams into daytime engineering**. They come

off the overnight shift and onto the work that needs them. At the same time, pull Tier-1 ops and IT back from outsourcing now that AI handles the volume in-house. The measure of success is unit cost per handled contact or incident down ~30%, with overnight headcount reduced and no degradation in service levels. If SLAs slip, the automation layer is under-resourced — reinvest before you cut again.

Stage 3 — Fold security in, CALEA-safe. Merge security's Tier-1/Tier-2 into the fusion center, with role-based access ring-fencing lawful-intercept activation to CALEA-designated personnel and board oversight preserved for NIS2. Keep specialized threat hunting and incident response as a daytime expert discipline. The gate here is the one that keeps a CISO in the room: mean-time-to-resolution improves with **no** increase in missed-security-alert rate, and the audit comes back clean.

Timeline, by operator scale. Because the work is AI-led, this moves faster than a traditional transformation — but “faster” still scales with size:

- **A small operator (or a Tier 3 served through a managed provider): ~90 days.** Few seats, one site, little political drag — the reference-customer sub-two-hour deployment shows how fast the technical part can move.
- **A Tier 2: about six months.** Enough scale to require real change management, still small enough to complete a full fusion in a single program.
- **A Tier 1: twelve to eighteen months.** At carrier scale — tens of thousands of ops staff, unions, multi-vendor estates, regulatory drag — you start delivering value in the first six months (Stage 1), but full consolidation is a phased twelve-to-eighteen-month program, and the deepest security fusion may run to the far end of that.

Through all three stages the spine is the same, and it is the message to lead with in every internal conversation: **this is people-first.** Your scarce, senior, hard-to-replace talent belongs on engineering — designing, building, solving — not on the 3 a.m. monitoring shift. Telecom is hit hardest of all industries by the retirement wave; the humane response and the economically correct response are, for once, the same response. Free your best people from the work a machine should do, and point them at the work only they can.

Fuse or fold

The macro conditions are not a phase. Flat revenue, expensive money, a retiring workforce, and 5G that didn't move the needle are the operating environment now, and they reward efficiency and monetization — not growth-at-any-cost. Running four separately-staffed round-the-clock teams, each with its own tools and its own truth, is a structure built for an era that has ended. It is expensive, it burns out the people you can least afford to lose, and it leaves you unable to hire your way out.

The Operations Fusion Center is the structure built for the era that has arrived. One team. Every signal. AI on the repetitive volume, humans on the moments that need them, your best engineers moved to daylight and to the work that matters. The industry is already converging on the fusion center for security. Telecom's opportunity is to take it one step further — to the customer, to the whole of operations — and to capture the efficiency the business is demanding while keeping the human touch the business will always require.

This is what has to happen. Macro conditions and the failure of the last technology cycle to deliver have made it not a strategy but a necessity. The only real choice left is how deliberately, and how soon, you move.

Fuse, or fold.

See it on your own numbers. Rapax deploys in five days and runs a 30-day evaluation against success criteria you write, with the \$25,000 evaluation fee credited against your license. Run the Fusion Savings Calculator at www.rapax.app, or book time directly at cal.com/shawn-ennis.

The Fusion Savings Model

The model behind the calculator. Every input ships with a sourced default; you override with your own figures.

Inputs (per function — CX, NOC, SOC, IT)

- **seats** — positions staffed at once.
- **coverage** — 24/7 or 8x5.
- **fte_per_seat** — default **5.0** for 24/7 (4.2 theoretical minimum times a ~1.2 relief factor); **1.0** for 8x5.
- **loaded_cost** — fully-loaded annual cost per FTE (default = base times **1.3**; base from BLS by role).
- **bpo_spend** — current annual outsourcing spend for this function.
- **tooling_spend** — annual ticketing/monitoring/inventory licensing attributable to this function.
- **deflection_rate** — expected AI automation of Tier-1 volume (benchmark **60–80%**; use a conservative starting value).
- **redeploy_share** — fraction of freed capacity redeployed to engineering vs. removed (a choice, not a benchmark).
- **platform_cost** — annual Rapax cost.
- **ramp_cost** — one-time implementation.

Core formulas

```
coverage_labor = seats * fte_per_seat * loaded_cost
current_run_rate = sum( coverage_labor + bpo_spend + tooling_spend ) # across 4 functions

gross_labor_saving = sum( seats * fte_per_seat * loaded_cost
                        * deflection_rate * consolidation_factor )
bpo_saving = sum( bpo_spend * insource_rate ) # insource_rate default 0.7
tooling_saving = sum( tooling_spend * consolidation_rate ) # default 0.5

gross_annual_saving = gross_labor_saving + bpo_saving + tooling_saving
net_annual_saving = gross_annual_saving - platform_cost
payback_months = ramp_cost / (net_annual_saving / 12)

# the split the CFO must choose:
budget_reduction = net_annual_saving * (1 - redeploy_share)
redeployed_to_eng = net_annual_saving * redeploy_share
```

consolidation_factor (default ~0.5–0.65) captures that one fused 24/7 front door replaces four, on top of raw deflection — this is the collapsed-coverage lever and the largest single driver. Present the result as a **range** (vary deflection_rate and consolidation_factor across their benchmark bands), never a single point, and always show the budget_reduction / redeployed_to_eng split so the number is honest about what it assumes.

Sources referenced in this paper include the U.S. Bureau of Labor Statistics; Gartner; Omdia; Analysys Mason; GlobalData; Dell'Oro; ISG; Times; Nemertes; company disclosures (SEC filings); and Rapax production data from its own operations. Cost figures are labeled sourced or modeled throughout; modeled figures are illustrative and depend on the assumptions shown.