Institute for Physical AI @ Bailey Military Institute · The Charlot Lab

FIELD OPERATIONS, ASSET MODELS AND THE COST OF A WRONG RECORD

Technical Report TR-2026-30
Market & Economic Review · v0.1
5 August 2026

FIELD OPERATIONS, ASSET MODELS AND THE COST OF A WRONG RECORD

Physical AI and Logistics Management for Telecommunications

Telecom's physical plant is large, poorly recorded and expensive to touch. This review prices what Physical AI can take on, grades every figure in the public record, and finds the deployed reality some distance behind the marketing.

Ronnie Gomez, Industrial Research Fellow · The Charlot Lab

Graded deployment, vendor and programme records

43 numbered sources, each carrying a verification grade

4 tables · verified / reported / self-published / modelled

Market & economic review · not investment advice

Telecommunications operators run one of the largest distributed physical estates in any industry and know its contents only approximately. CTIA counts 447,605 US cell sites at the end of 2024 [6]. WIA's 2025 report, modelled by iGR, estimates 254,850 macrocells and 158,500 purpose-built towers [4], and neither body publishes a methodology that reconciles the two. The buried-plant record is better documented, and what it documents is worse: the Common Ground Alliance logged 196,977 damages to buried utilities in 2024, telecom accounting for close to half, with about one damage in five traced to a facility that was unmarked or marked in the wrong place [31]. This review prices what Physical AI can take on against that background. Labour is 72% of the labour-plus-materials line for underground fiber and 67.9% of median total deployment cost on the Fiber Broadband Association's own figures [1,42], and WIA puts US cellular opex at $54.7bn against $10.2bn of capacity capex [4]. Structural tower inspection, the most heavily marketed application, is worth on the order of $21-36M a year nationally at

vendor-quoted prices, under a tenth of a percent of that opex [42]. The deployed record also runs behind the marketing: in the FAA's flagship BVLOS programme, 98% of flights still carried a visual observer [23]. The report closes with a measurement programme an operator can run before committing capital.

1. Headline numbers

The table below collects the figures this review returns to. Grades follow the Institute's convention: verified means independently replicated or drawn from a regulator, standards body or peer-reviewed source; reported means a named analyst house or operator stated it without independent confirmation; self-published means the number comes from a party selling the thing being measured; modelled means it was computed rather than observed. Where a figure is the author's arithmetic on a cited primary, it carries index [42] and is graded modelled throughout the paper.

Two of these entries deserve flagging at the outset because they reframe the rest. The first is the ratio of recurring operating cost to expansion capital in US wireless, which WIA's own numbers put at 5.4 to 1 [4,42]. The second is the share of FAA showcase-programme flights that were genuinely uncrewed at the site, which was under 2% [23]. The first says the prize is in operations rather than buildout. The second says the mechanism most often proposed for capturing it is not yet legally available at scale.

Key number	Value	Grade	Source
US cell sites, year-end 2024	447,605 total, of which 166,264 small cells	reported	CTIA [6]
US wireless estate, 2025	254,850 macrocells; 158,500 purpose-built towers; 198,100 outdoor small cells; 830,350 indoor nodes	modelled (iGR)	WIA / iGR [4]
US cellular opex vs capacity capex, 2025	$54.7bn opex against $10.2bn capex; ratio 5.4x	modelled (iGR), ratio modelled	WIA / iGR [4,42]
Implied cost per site per year	$214,600 per macrocell, or $42,600	modelled	[4,42]

	per node across all site types		
Fiber deployment cost, 2025 median	$18.00/ft underground, $8.00/ft aerial; about $95,000 and $42,000 per mile	reported (per-mile modelled)	FBA / Cartesian [1,42]
Stated labour share of fiber cost	72% underground, 64% aerial	reported	FBA / Cartesian [1]
Labour as share of the median total cost	67.9% underground, 56.2% aerial	modelled	[1,42]
Cost direction	92% of respondents reported 2025 increases; 88% expect more in 2026	reported	FBA / Cartesian [1]
Global telecom labour opex	$258.1bn (MTN Consulting, 140 telcos) against about $300bn (Omdia)	reported	[19,20]
Global telecom capex, annualised 2Q25	$292.9bn, down 2.6% year on year, lowest 12-month total since 2011	reported	MTN Consulting [19]
Field-service first-visit resolution	87% median; engineer utilisation 83%, 90.6% among top performers	reported	TSIA [15]
Truck roll cost	$150–500 direct; about $1,000 fully loaded	reported (weakly sourced)	TSIA via vendor literature [14]
Buried-utility damages, 2024	196,977 reported; telecom close to half; 20.5% from locate or marking failure	reported	CGA DIRT [31]
	Fewer than 763 of more than 44,000 in	verified	

BVLOS flights flown without a visual observer	FAA BEYOND (about 2%); zero in PSP/IPA		US DOT OIG [23]
Infrastructure inspection share of BEYOND operations	About 1% of more than 60,000 flights, roughly 600 flights over four years	verified	US DOT OIG [23]
Tower structural inspections implied by TIA-222	About 32,000–53,000 per year in the US	verified standard, modelled on a modelled base	[7,42]
National value of drone structural inspection at vendor prices	$21–36M per year, 0.04–0.07% of US cellular opex	modelled	[4,7,8,42]
Vendor digital-twin O&M saving claim	$5,700 falling to $2,200 per tower per year	modelled (vendor)	Bentley [29]

Table 1. Reference figures for the US and global telecom physical estate, with verification grade.

2. Scope, definitions and method

Physical AI in this review means machines that sense, move, manipulate or navigate in the field and decide onboard what to do next. That covers uncrewed aircraft flying tower orbits, ground vehicles working inside exchanges and data halls, robotic boring and trenching equipment, and the perception and planning software that turns sensor output into a usable model of the plant. It also covers the diagnosis and dispatch systems that decide whether anything gets sent at all, because on the current evidence those carry more measured value than any embodied system in telecom.

Logistics management means the allocation and movement of physical resources against that plant: crew routing and scheduling, spares and consumables, work sequencing, and restoration after storm or excavation damage.

The review excludes network-layer autonomy that never produces a physical action. Self-optimising RAN, traffic steering and energy-saving sleep modes are large subjects, well covered elsewhere, and they enter here only where the output is a

decision to dispatch or not to dispatch. Customer-premises equipment and consumer self-install are also out of scope except where they bear on truck-roll avoidance.

The fellowship framing comes from Ronnie Gomez's project statement [43], which names four capabilities: locate and inspect assets, route field crews, track inventory, and coordinate the people and machines that service the network. The statement also names the scientific problem correctly, which is giving those systems a reliable model of the physical plant and the judgment to act on it when conditions get messy. Sections 3 and 9 take that problem as the organising question. Sections 5 through 8 test each of the four capabilities against what the public record can actually support. One of the four, inventory, turns out to have almost no public evidence at all, which is itself a finding and is recorded in section 11.

Method. Two parallel evidence tracks were compiled on 4-5 August 2026, one on economics and one on technology, and merged into 43 sources. Regulator, standards-body and audit documents were preferred where they exist, which in this domain means the US DOT Office of Inspector General [23], OSHA [11] and ANSI/TIA-222 [7]. Trade-association datasets are used where nothing better exists and are marked as member-supplied. Where a figure could only be found in material published by a party selling the product, it is graded self-published and the seller is named in the same sentence.

Limits worth stating up front. bls.gov blocked automated retrieval during compilation, so reference [3] rests on search-index extracts of the Occupational Outlook Handbook rather than a direct read of the primary page. The TSIA benchmark data in [15] sits behind membership and is quoted from secondary description. The peer-reviewed review at [41] was truncated in the compiled evidence record and its full citation needs completing before publication; it is used here only for the qualitative claim that contact servicing of towers remains at feasibility stage. Nothing in this report is investment advice.

One tie-break is needed because the two middle grades overlap. Where a party is both the operator of a programme and the seller of the thing being measured, or is stating a result about its own operations, the claim is graded self-published regardless of which outlet relayed it. Trade-association figures are marked member-supplied. Trade-press repetition of a company statement is not independent confirmation and does not raise a grade.

3. The estate and its counts

Any autonomous system that acts on telecom plant needs to know what is there. The industry does not agree on the answer even in aggregate. CTIA's 2025 survey, drawn from its own members, reports 447,605 cell sites in operation at year-end 2024, of which 166,264 were small cells, up 6% from 156,787 the year before [6]. WIA's 2025 report, modelled by iGR, gives 254,850 macrocell sites, 158,500 purpose-built towers, 198,100 outdoor small cells and 830,350 indoor small cell nodes [4]. WIA's indoor node count on its own exceeds CTIA's entire site total. The two counts cannot be reconciled without underlying definitions that neither body publishes, and both are association surveys of member-supplied figures with no independent audit.

This matters more than a footnote about survey methodology. CTIA's 447,605 and WIA's 254,850 macrocells differ by 1.76x, and against WIA's 1,283,300 all-node total the ratio is 2.9x; the two also differ in reference date and in scope. That spread at national level is the top of a hierarchy of registry error that runs all the way down to individual assets. In a Bentley-sponsored validation of the OpenTower iQ digital twin, one of four towers examined had a RAD centre documented at 180 ft that the twin measured at 175 ft [28]. A five-foot error on a parameter that drives structural loading and RF design, in the one case of four the study details, is the concrete version of the same problem. The study does not establish that the other three sites were correct on that parameter.

Below ground the evidence is much better and the picture is worse. The Common Ground Alliance's 2024 DIRT Report, published 28 August 2025, analysed 196,977 unique damage reports to buried utilities across North America, with telecommunications facilities accounting for close to half of reported damages [31]. This is the largest dataset bearing on the fellow's problem and the only one that measures the cost of a wrong plant model directly, in units of broken infrastructure.

The root-cause breakdown is where the operational signal sits. CGA attributes 24.5% of damages to failure to notify 811, 16.1% to an excavator failing to maintain clearance, 11.9% to a facility not being marked at all, 8.6% to inaccurate marking, and 6.8% to improper excavation [31]. Two of those categories, together 20.5% of damages, are locator error, which works out to roughly 40,400 damages a year traceable to the buried-asset model being wrong or unfindable in the field [31,42]. If telecom's share of those follows its share of damages overall, the telecom-specific

figure is on the order of 20,000 a year, though CGA does not publish the cross-tabulation that would confirm it [42].

The five root causes listed above are a subset of the top ten to which CGA attributes 85% of reported damages, and they sum to 67.9% [31,42]. CGA itself treats DIRT as an undercount, since submission is voluntary and the reporting population shifts year to year, and the CGA Index moving from 94.0 to 96.7 partly reflects who reported rather than what happened.

There is a second, schedule-shaped cost in the same report. Across eight 811 centres, excavators faced an average 38% chance of being unable to start work on time because locate responses were incomplete [31]. That figure covers eight centres and is not nationally representative, but it puts a number on what an untrusted registry does to a construction programme before anything is dug at all.

4. Is field labour scarce? The sources disagree

The standard case for automating telecom field work opens with a labour shortage. The numbers behind that case do not agree with each other, and the disagreement is not marginal.

On one side, Continuum Capital, commissioned by the Fiber Broadband Association and the Power & Communication Contractors Association and presented at Fiber Connect in July 2024, put the requirement at about 28,000 additional broadband construction workers and 30,000 additional broadband technicians to absorb planned federal and state funding, plus roughly 119,000 more over ten years to cover retirement and attrition, with the shortfall pushing projects two to three years out [2]. It is a forecast, commissioned by two trade associations whose members benefit from workforce funding, published before both the BEAD restructuring and the 2025 contraction in wireless construction. The Communications Workers of America has publicly called US telecom workforce shortage claims misinformation [2].

On the other side, WIA's own 2025 data shows the US wireless infrastructure workforce falling to 342,350 full-time equivalents from 368,750 in 2024, a decline of 26,400 jobs or 7.2% [4]. Wireless Estimator, tracking against its own twenty-year record, reported tower construction hiring at a two-decade low in Q4 2025, with help-wanted postings down about 60% against recent years while jobseeker registrations rose more than 300% [5]. That source is a single proprietary job board

reporting on its own commercial base, and postings volume is a noisy proxy, but the direction is unambiguous and it is the opposite of the fiber trade bodies' direction in the same period.

These are not strictly the same labour pool. Fiber construction and wireless tower work draw on overlapping but distinct crews, and a fiber shortage can coexist with wireless slack. The honest reading is that the shortage is regional, segment-specific and unresolved in the public data, which means an automation business case built on wage scarcity is building on a number nobody has established.

The BLS projection points somewhere more useful. Employment of telecommunications equipment installers and repairers, excluding line installers, is projected to decline 3% from 2024 to 2034, while about 23,200 openings a year are projected, all of them arising from workers transferring out of the occupation or leaving the labour force [3]. Demand for the occupation is shrinking and the replacement requirement is nonetheless large. What walks out with those workers is undocumented site knowledge: which cabinet the spare fibre is coiled in, which gate code works, which pole is the one that was replaced after the 2019 storm. That is a knowledge-capture problem with a machine-readable answer, and it is a better argument for asset digital twins than any wage forecast.

Safety is the other standing argument for removing people from towers, and it rests on a rate nobody agrees on. OSHA's own count, which has not been updated past 2016, records 13 communication tower fatalities in 2013, 12 in 2014, 3 in 2015 and 6 in 2016 [11]. The rate dispute is sharper. CWA derived 19.3 deaths per 100,000 workers from a five-year average of 5.6 deaths over a 29,000-worker denominator; safety experts using a workforce denominator of about 37,000 and a fatality database put 2021 at 10.8 per 100,000 on 4 deaths and 2022 at 8.4 on 3 deaths [12]. Decomposing the 1.8x gap, about 1.4x of it comes from the numerator, a five-year average against a single year, and about 1.28x from the denominator, 37,000 against 29,000 [12,42]. Both parties are interested; CWA's survey had 103 respondents and Wireless Estimator's database is proprietary and internally inconsistent on its own date range. There is no BLS occupational code isolating tower technicians, so no neutral rate exists. NATE's framing is the more actionable one: technicians typically ascend and descend six to ten times a day, and eliminating two or three of those climbs is argued to reduce fatigue-driven risk [13]. That is a claim about removing some climbs from a day rather than removing the crew, and it matches what is actually deployed.

5. The cost stack of a fiber mile

The Fiber Broadband Association's third annual deployment cost report, produced by Cartesian and published in January 2026, is the best public per-unit benchmark in this area and the denominator any automation claim has to beat. Median reported cost to deploy fiber in the US in 2025 was $18.00 per foot underground and $8.00 per foot aerial, roughly $95,000 and $42,000 per mile [1,42]. Underground rose 3% year on year and aerial 14%. The sample covers projects in 38 states, gathered by phone interview and online survey in September and October 2025, self-reported by operators and contractors, with sample size undisclosed in the executive summary and each year's panel independent, so the year-on-year comparison is not like-for-like.

Cartesian reports labour at 72% of underground deployment cost and 64% of aerial, with median labour of $12.23/ft and $4.50/ft against median materials of $4.65/ft and $2.50/ft [1]. Those shares reproduce exactly as labour divided by labour plus materials: 72.5% and 64.3% [42]. Against the median total cost figures they are 67.9% and 56.2% [42]. The gap means about $1.12/ft underground and $1.00/ft aerial, or 6.2% and 12.5% of the reported total, sits outside the labour-and-materials frame in permitting, make-ready, engineering and surface restoration [42]. Medians of different quantities are not additive, so this decomposition is indicative rather than exact, but the exact match of both stated shares to the labour-plus-materials denominator makes the inference reasonably firm. Anyone modelling a labour-productivity saving against total project cost should use the lower share.

The most useful calibration point in the report is the internal-versus-contractor comparison. Using internal crews rather than contractors cut labour cost 13% underground, $11.13/ft against $12.73/ft, and 11% aerial, $4.00/ft against $4.50/ft [1]. Cartesian notes that internal and external labour figures may carry different context by respondent, and some internal numbers may exclude overhead that contractor invoices include. Taken at face value, that delta is worth $8,448 per underground mile and $2,640 per aerial mile, which is 8.9% and 6.2% of total deployment cost [42]. Vertical integration is the best-documented organisational lever available, and it moves total cost by single digits. An autonomous system claiming materially more than that is claiming something structurally different from better scheduling, and should be asked to say what.

Method selection is the larger lever and the one most amenable to an automated planner working from terrain and subsurface data. Trenching carries the highest

cost of the common methods, 60% above plowing and 6% above directional boring on median costs, and direct burial runs about 40% cheaper than laying in new conduit [1]. Moving a mile from trenching to plowing where ground conditions permit is worth about 37.5% of that method's cost [42], which is several times any software-driven productivity claim in the field-service literature. Cartesian notes that direct burial offers less flexibility and potentially higher maintenance cost, so the 40% is a first-cost figure with no published lifecycle counterpart.

Deutsche Telekom's own figures make the same point at a different scale: about €85 per metre for traditional underground deployment, about €65 for reduced-depth less-invasive methods, and about €10 per metre overground [36]. Going aerial is roughly an eight-fold saving, larger than anything a robot has demonstrated. These are the operator's own numbers with no published methodology, and German civils costs are dominated by local reinstatement and permitting rules that do not transfer. For orientation, the US median underground figure converts to about $59 per metre [42], the same order once currency is set aside.

The cost curve is moving the wrong way and the respondents name the reason. 92% reported deployment cost increases in 2025 and 88% expect further increases in 2026, attributing them primarily to labour rates, then materials, permitting and make-ready [1]. The expectation data has one year of calibration behind it: 80% expected a 2025 increase and 92% then reported one, so this panel under-forecasts rather than over-forecasts, on a small self-selected sample.

Component	Underground (median)	Aerial (median)	Note
Total deployment cost	$18.00/ft, about $95,040/mi	$8.00/ft, about $42,240/mi	Per-mile conversion is the author's [42]
Labour	$12.23/ft	$4.50/ft	Fully loaded as each respondent defines it [1]
Materials	$4.65/ft	$2.50/ft	Driven by tariffs and commodity prices [1]
Labour plus materials	$16.88/ft	$7.00/ft	Sum of medians; medians are not additive [42]

Residual against the total median	$1.12/ft (6.2%)	$1.00/ft (12.5%)	Permitting, make-ready, engineering, restoration [42]
Labour share as published	72%	64%	Reproduces as labour ÷ (labour + materials) [1,42]
Labour ÷ total median cost	67.9%	56.2%	Author's arithmetic [42]
Internal vs contractor labour	$11.13 vs $12.73/ft, down 13%	$4.00 vs $4.50/ft, down 11%	Definitions may differ by respondent [1]
Value of that delta per mile	$8,448 (8.9% of total)	$2,640 (6.2% of total)	Author's arithmetic [42]
Value of a 10% labour productivity gain	$6,457/mi	$2,376/mi	155 underground miles to save $1M [42]
Method spread	Trenching 60% above plowing, 6% above boring	Not applicable	First cost on median figures [1]
Conduit vs direct burial	Direct burial about 40% cheaper	Not applicable	First cost only; no lifecycle comparison published [1]

Table 2. The 2025 US fiber deployment cost stack, with author's derivations marked. Source: FBA/Cartesian [1]; derivations [42].

6. Recurring field operations

WIA and iGR put 2025 US cellular network operating expenditure at about $54.7bn against more than $10.2bn of capacity-and-coverage capex [4]. Recurring operations run at 5.4 times expansion capital on those figures [42]. The $10.2bn is capacity and coverage only and is not comparable to CTIA's roughly $30bn total wireless capex for 2024, which uses a wider definition [6]. Even allowing for that, the ordering is the fellow's thesis stated in the industry's own numbers: the money is in keeping the network running.

Per-site denominators are where this gets slippery. Dividing $54.7bn by WIA's 254,850 macrocells gives about $214,600 per macro site per year, but the opex denominator also covers 198,100 outdoor small cells and 830,350 indoor nodes; spreading it across all 1,283,300 nodes gives about $42,600 [4,42]. The honest bracket for cost per site per year is therefore $43,000 to $215,000 depending entirely on which sites you think the opex belongs to, and WIA does not publish the split. Any business case denominated in cost per site should state which end of that range it is using.

At global scale the two most-cited houses disagree on the single largest line. MTN Consulting's Telco Market Tracker for 2Q25, published December 2025, covers 140 operators and puts annualised capex at $292.9bn, down 2.6% year on year and described as the lowest twelve-month total since 2011 and the lowest capital intensity since 2013. Labour costs came to $258.1bn, down 1.3%, across headcount of 4.36 million, with labour at 21.7% of opex against 22.2% a year earlier and EBITDA margin at 33.8% [19]. Omdia's Global Telecoms Opex Tracker puts labour opex at about $300bn, falling 2.2%, within total global opex of $1.63 trillion against revenue of $2.01 trillion [20]. The two figures are 16% apart [42], but they do not cover the same universe: MTN's $258.1bn spans 140 named telcos and Omdia's is a global total, so an unknown part of the gap is population coverage rather than disagreement about the same quantity. MTN's own annualised revenue base of about $1.8 trillion sits between Deloitte's $1.55tn and Omdia's $2.01tn. Omdia also attributes part of the 2024 decline to divestments rather than efficiency, and notes that disclosures of quantifiable AI savings remain scarce. A later MTN tracker is published; 2Q25 is used throughout for consistency with the labour split, and the vintage is stated here rather than left implicit.

MTN's figures yield an average fully loaded labour cost of about $59,200 per telecom employee per year [19,42]. That is the most useful single anchor available, and it is also where the public data stops being useful, because no source isolates field-operations labour from total telecom labour. Headcount of 4.36 million includes retail, network engineering, IT, customer care and back office. The field-crew subset is the number this entire subject depends on and nobody publishes it.

Capital intensity is similarly contested. Dell'Oro's 2026 forecast, covering about fifty service providers representing roughly 80% of global capex, has capex flat in nominal terms in 2025, declining 2% in 2026, then growing at about 1% CAGR to 2030, with capex-to-revenue approaching 14% by 2029 and wireless capital intensity around 11%, down seven points from the 5G peak [18]. Dell'Oro's own

September 2025 release gave 15% and 12-13% for the same horizon, so the house revised itself by roughly a point of total intensity and one to two points on wireless within twelve months. Deloitte's 2026 outlook puts capex intensity around 18% against global revenue of roughly $1.55 trillion citing IDC [21]. Three named houses span 14% to 18%, on revenue bases that themselves differ by 1.3x between Deloitte's $1.55tn and Omdia's $2.01tn [42]. The size of Dell'Oro's own year-on-year revision is comparable to the total effect a successful field-automation programme would be expected to have on capital intensity, which means claims to detect automation savings in aggregate industry statistics are claiming to see a signal smaller than the forecasting noise.

Carrier guidance shows divergence rather than a trend. AT&T guided to roughly $22bn of capital investment in 2025 and $23-24bn in 2026; Verizon to $17bn in 2025 falling to $16-16.5bn in 2026; T-Mobile to $10bn in both years [22]. In aggregate the three are close to flat while pulling in different directions.

Field-service benchmarking gives the realistic headroom. TSIA puts the industry median at 87% of on-site incidents resolved in a single visit, with engineer utilisation at 83% industry-wide and 90.6% among top performers [15]. Roughly one visit in eight fails first time, and about a sixth of technician time is already non-productive. Closing the utilisation gap to top-performer level is worth 9.2% more productive hours from the same headcount, with no new technology at all [42]. That is the bar an autonomous scheduling system has to clear before it has done anything.

The number everything in this field is denominated in is also the weakest one in it. A truck roll is commonly costed at $150-500 direct, rising to roughly $1,000 fully loaded, with the $1,000 attributed to TSIA [14]. No primary TSIA publication stating $1,000 could be located; every citation runs through a vendor selling truck-roll avoidance. The range spans nearly sevenfold and mixes residential broadband installs with enterprise equipment service. Operators building a case on cost per site should use their own audited figure and treat the published range as an order-of-magnitude anchor only.

Against that, the operator-attributed results are modest and consistent with each other. Vodafone reduced technician dispatch rate by 10%, raised first-contact resolution five points and raised NPS fifteen points using remote visual assistance, published by the supplier TechSee with no date, market or baseline [16]. Vodafone's Field Technician Assist, built by Celfocus on Google Cloud, delivers top-three recommended actions to field engineers and is credited with a 28% reduction in

repeated field visits [32]. Applying that to TSIA's 13% repeat rate removes about 3.6% of total visits [15,32,42]. Comcast reports a machine learning system predicting with about 90% accuracy whether a technician needs to be dispatched, credited with saving tens of millions of dollars in avoided truck rolls [17,34]. Comcast does not size the saving further and no denominator is applied to it here: the truck rolls sit in Connectivity & Platforms, and dividing a segment-level saving by total corporate revenue, which includes NBCUniversal and Sky, would measure it against the wrong base. Accuracy without a precision and recall split is close to meaningless for an imbalanced dispatch decision where a false negative leaves a customer broken and a false positive costs a van. Elsewhere in the same vendor literature, claims of 60% truck-roll reduction appear with no operator attribution at all.

7. What is actually deployed

The largest asset-twin programme in the public record is Verizon's. Verizon states that it has drone-derived twins across tens of thousands of cell sites and, from early 2025, automatic capture and digitisation of every new build and modification, with drone capture cutting some disaster-response site assessments from several hours to about fifteen minutes [25]. The software partner is vHive, which raised a $25M Series B in April 2022 led by PSG and reports generating tens of thousands of twins in a year across telecom, construction and cranes, with Verizon and Deutsche Telekom named as customers [26]. Deutsche Telekom's Telekom Innovation Pool participated in that round, so DT is an investor in the supplier as well as a named customer of it. Consilience Analytics states it has processed over 35,000 twins since 2019 [30]. That count is global, spans sectors and is cumulative across repeat captures, so it cannot be set against the US tower estate as a coverage figure and no such ratio is computed here. No independent confirmation of scale, cadence or accuracy has been published for either programme.

The mechanism inside Verizon's programme is worth separating from the hardware. Capture on every change is a process commitment that keeps a plant model from drifting; the drone is the instrument that makes it affordable. Site count is given as a range rather than a figure, and no accuracy metric, cost per twin or refresh cadence for the legacy estate has been published. American Tower and Crown Castle both build drone-derived twins, and American Tower reports being rated industry leader in quality of asset information in its own 2025 customer survey,

which has no published methodology, sample or comparison set [10]. Neither towerco discloses how many structures have been scanned or at what accuracy.

On accuracy the fullest public disclosure is Bentley's sponsored validation of OpenTower iQ, which found automated equipment detection and dimension measurement 95% accurate and tilt and azimuth accurate within a 5% error margin, with tower height on one example reading 72 ft 3 in against 73 ft in the design documents, a 1.03% difference [28,42]. The sample was four towers drawn from a group of fifty. An accuracy rate from four observations carries a binomial interval spanning most of the usable range, so the 95% is recorded here as the vendor's characterisation and not as a measurement. The comparison baseline was best-existing data, meaning RF data sheets, old structural analyses and CAD drawings ranging from under a year to several years old. That is the crux: the study measures agreement with the records, and the reliability of the records is the thing in question. The paper does not state whether the 5% is of reading or of full scale.

Bentley's economics are a model rather than a measurement. The company estimates manual tower operations and maintenance at about $5,700 per tower against $2,200 with OpenTower iQ, a $3,500 annual saving, alongside 50% reporting time savings, a 15% annual opex reduction and 42% lower ten-year TCO per tower [29]. The per-tower figures imply a 61% cut in the stated O&M line, which the same document does not reconcile with its own 15% and 42% claims; the three percentages are computed over different denominators that are never named [42]. No operator has published a matched-cohort before-and-after. The model excludes the cost of maintaining twin currency, re-capture after each site modification, and pilot and mobilisation cost, all of which fall on the operator. Applied to WIA's 158,500 purpose-built towers, the $3,500 would come to about $555M a year across the entire US tower estate, roughly 1% of WIA's $54.7bn cellular opex, though the two are different denominators and the comparison is only for scale [4,29,42].

Aerial structural inspection has the loudest economics and the weakest sourcing. Providers quote $350-600 per tower against $800-1,500 for a climbing crew, with drone crews completing four to six towers a day against one to two for a two-person climbing crew, and 30-50% overall cost savings [8]. Every figure traces to a company selling the service, and no independent, operator-audited or peer-reviewed comparison could be found. The comparison is also not scope-matched, since a drone pass performs no torque checks, connector work, grounding continuity testing or hands-on remediation. On the quoted midpoints the implied saving is

59%, and the full spread of the quoted ranges runs from 25% to 77% [42], so the marketed 30-50% is itself a selection from a much wider band.

ANSI/TIA-222 recommends major structural inspection every three years for guyed masts and every five years for self-supporting structures, plus after severe weather, and TIA's TR-14 committee states explicitly that this is a recommendation rather than a mandate [7]. Applied to 158,500 purpose-built towers, that implies roughly 32,000 to 53,000 structural inspections a year depending on the guyed and self-support mix, which is not published [7,42]. At the quoted midpoint saving of $675 per tower, the entire national value of drone structural inspection is $21-36M a year, which is 0.04-0.07% of US cellular opex [4,8,42]. AT&T announced drone tower inspections in 2016 across roughly 65,000 cell sites where it has equipment, as characterised in trade press at the time, and stated that the programme would break even during the initial ramp-up period; an AT&T official declined to quantify the saving beyond that [9]. No later cost figure from AT&T has been located in public material. The absence is not itself evidence of the size of the effect.

Physical work on the plant is thinner still. Openreach is trialling MicroBORE, a trenchless robotic drilling system developed over five years by BT Group, Synthotech and the University of Surrey from lunar and Martian drilling research, to bore narrow channels under pavements and driveways for FTTP drops; early testing at BT's Suffolk research centre confirmed operational viability and broader deployment is under evaluation [35]. No bore rate, cost per drop, failure rate or comparison against manual moling has been published, and Openreach's £15bn programme to reach 25 million premises by December 2026, about £600 per premises passed, is being met without it [35,42]. Built Robotics retrofits standard excavators for autonomous trenching in about a day with contractors reporting 20-30% savings, but its volume business is autonomous solar pile driving, including a nationwide agreement with Blattner announced in September 2025 [37]. The best-funded autonomous trenching company found product-market fit in open, repetitive, greenfield geometry. Telecom conduit is short-run, congested with existing utilities and constrained by surface restoration, which is close to the opposite.

Ground robots inside telecom are mostly not doing telecom work. AT&T fields Ghost Robotics quadrupeds over FirstNet using its Geocast service for beyond-line-of-sight command and control, with documented use in public safety and defence, including Air Force flight-line and perimeter patrol at Tyndall AFB, rather than inspection of AT&T's own plant [38]. That is a connectivity product sold to others.

NTT DATA began rolling out data centre facility inspection robots across its Japanese sites from April 2023 and moved to commercialise the offering [39], which is the clearest at-scale ground-robot inspection programme inside telecom-owned facilities and sits in controlled indoor environments rather than central offices or outside plant. The operator-led embodied deployments are happening in Asia: at MWC Shanghai in June 2026 China Telecom announced a Robotics-as-a-Service model with AgiBot for aviation and logistics, China Unicom described robot deployment for hazard inspection in chemical plants, and State Grid has announced plans for roughly 8,500 robots, about 5,000 quadrupeds, 3,000 dual-arm wheeled units and 500 humanoids, across grid inspection and maintenance [40]. Those are announced programme figures from state-linked entities with no independent verification of units in service, uptime or task success. If the model works it is a services line for carriers, which is a different business thesis from internal cost reduction. Peer-reviewed work on robots that physically service telecom towers remains at feasibility stage [41].

Programme	Operator / vendor	What is documented	Grade
Asset twins at scale	Verizon with vHive	Tens of thousands of sites; automatic capture on every build or modification from early 2025; disaster assessment cut from hours to about 15 minutes	self-published [25]
Twin production volume	Consilience Analytics	35,000+ twins processed since 2019; includes repeat captures	self-published [30]
Twin accuracy	Bentley OpenTower iQ	95% equipment detection; tilt and azimuth within 5%; baseline was existing records, not climb-verified truth; n=4 towers	self-published [28]
Twin economics	Bentley	$5,700 falling to $2,200 per tower per year; excludes recapture, mobilisation and currency maintenance	modelled [29]

Towerco twin programmes	American Tower, Crown Castle	Drone-derived twins; AMT rated leader in its own customer survey; no scan counts or accuracy published	self-published [10]
Aerial structural inspection	Five named service providers	$350–600 per tower against $800–1,500 for a climbing crew; scope not matched	self-published [8]
Carrier drone inspection	AT&T	In use since 2016–17 across roughly 65,000 towers; no cost figure ever published	self-published [9]
Dispatch prediction	Comcast	About 90% accuracy on whether a visit is needed; tens of millions saved; no precision or recall split	self-published [17,34]
Storm restoration	Comcast	Alarm clustering plus tooling dispatch; trialled in the 2024 hurricane season, nationwide from September 2025; 50% more effective, undefined	self-published [33]
Technician decision support	Vodafone with Celfocus	Top-three recommended actions from intervention history; 28% fewer repeated field visits	self-published [32]
Remote visual assistance	Vodafone with TechSee	10% lower dispatch rate; first-contact resolution up five points	self-published [16]
Robotic fibre boring	Openreach with Synthotech and University of Surrey	MicroBORE trenchless drilling; trial at BT Suffolk research centre; deployment under evaluation	self-published [35]
Autonomous trenching	Built Robotics	Excavator retrofit in about a day; 20–30% contractor-reported	self-published [37]

		saving; volume business is solar piling	
Legged ground robots	AT&T with Ghost Robotics	Sold as a FirstNet connectivity product for public safety and defence; no documented use on AT&T plant	self-published [38]
Indoor inspection robots	NTT DATA	Data centre facility inspection across Japanese sites from April 2023; no fleet count or measured saving	self-published [39]
Operator robotics-as-a-service	China Telecom, China Unicom, State Grid	RaaS with AgiBot; about 8,500 robots planned at State Grid	self-published [40]
Contact servicing of towers	Academic literature	Feasibility stage; no fielded system	reported (citation incomplete) [41]

Table 3. Deployment ledger: what is documented, and at what verification grade.

8. The regulatory gate on uncrewed inspection

The US DOT Office of Inspector General audited the FAA's beyond-visual-line-of-sight efforts and published Report AV2025034 on 30 June 2025. It is the only independent audit of BVLOS operational practice located by this review, and it does not support the picture the vendor literature paints.

The number vendors cite is that FAA approvals for BVLOS operations rose from 1,229 in 2020 to 26,870 in 2023, across small-UAS rule waivers, air carrier operating certificates and regulatory exemptions [23]. That is an approvals count, not a measure of autonomous flying. An approval can cover a single flight or thousands, the categories are not comparable, and the count says nothing about mission type or autonomy level.

The number that matters is different. Of more than 44,000 BVLOS flights flown under FAA's flagship BEYOND programme through October 2024, fewer than 763, about 2%, were flown without a visual observer. In the parallel PSP and IPA programmes, no operator flew a single flight without a visual observer [23]. The

economic case for drone tower inspection rests on removing the human from the site. In the FAA's own showcase programme, the human was standing there for 98% of flights.

Mission mix is equally instructive. Out of more than 60,000 BEYOND operational flights, 97% were package delivery, about 1% were infrastructure inspections and under 1% were public safety [23]. Infrastructure inspection is one of BEYOND's three named mission types and it produced roughly 600 flights across four years. The regulatory learning curve that will eventually govern telecom tower BVLOS was built almost entirely on parcel delivery, whose flight profile of long straight legs and low hover with no close approach to a radiating structure has little in common with a tower orbit. An FAA manager conceded to the OIG that the programme's KPI design was more tailored to package delivery.

The waiver base is thin and concentrated. FAA had issued 190 BVLOS waivers to 134 distinct operators as of October 2024, up from 6 in 2020 and 122 in 2023. Two BEYOND lead participants accounted for 95% of all BVLOS operations, and one participant alone flew over 60% of the no-observer flights [23]. Waiver counts are not broken out by industry, so how many of the 134 are telecom or tower-sector operators is not public. An operator planning a national BVLOS inspection programme is joining a queue that has cleared roughly 134 organisations in five years, not entering a mature market.

Two caveats cut in the operator's favour. The OIG's figures cover only the partnership programmes, so commercial waiver holders outside them are not counted, and the OIG separately found FAA's data validation process vulnerable to errors, having mis-categorised at least one operator's observer status. More importantly, telecom operators largely inspect under ordinary Part 107 within visual line of sight, so their flight volume never enters this dataset at all. What the dataset establishes is not that telecom drone inspection is rare. It is that the specific capability the savings depend on, flying a site without a crew present, is rare everywhere.

The gate is Part 108. FAA's BVLOS rule was proposed on 7 August 2025, drew roughly 3,000-3,100 comments, and the final rule reached OIRA review on 10 July 2026, with publication expected late 2026 or early 2027 and a six to twelve month transition after that [24]. None of the trade sources reporting this has an FAA commitment date, and OIRA timelines slip routinely. The rule's actual detect-and-avoid and operator-certification requirements are not public, so the cost of compliance is unknown and could be large enough to change the per-site economics

on its own. Until it lands, per-site waivers and visual observers remain the operating reality, and the labour saving that justifies the capital is not available.

9. Requirements for an actionable plant model

The fellow's project statement identifies the scientific problem as giving embodied systems a reliable model of the physical plant and the judgment to act on it when conditions get messy [43]. The evidence assembled here decomposes that into five problems, none of which is solved and each of which is measurable.

Ground truth. The Bentley validation compared the digital twin against best-existing data, meaning the RF data sheets, structural analyses and CAD drawings on file [28]. The twin agreed with the records 95% of the time on equipment detection. What the study cannot say is whether the twin or the record was right in the 5% where they disagreed, and the one case it details, a RAD centre documented at 180 ft and measured at 175 ft, is presented without saying which was subsequently confirmed. A validation that treats the registry as truth cannot detect the registry error that is the whole reason to build the twin. The fix is not exotic: a climb-verified or total-station-verified sample of thirty or more structures, stratified by tower type and age, with per-parameter agreement reported and disagreements adjudicated by physical measurement. No such study exists in public.

Navigation at a radiating structure. vHive states that electromagnetic interference around telecom towers disrupts RTK positioning badly enough to force manual flights and repeat site visits, and announced in November 2025 that it had restored autonomous capture on the DJI Matrice 4E in high-interference zones with a roughly five-minute automated 360-degree survey [27]. No positioning-error figures, no degradation curve against transmit power or frequency, and no mechanism description accompany the claim, and the sole-provider assertion is not independently checkable. The physical situation is nonetheless real and specific: the asset being modelled is a high-power RF emitter and the sensor modelling it depends on a weak GNSS signal. The binding constraint at a cell site is state estimation under partial GNSS denial rather than perception, which points toward onboard visual-inertial estimation with the structure itself as the reference frame rather than toward better GNSS. That work belongs on the device, because the backhaul at a site is least reliable during exactly the storm conditions when the inspection is most needed.

Accuracy budgets tied to tasks. A twin quoted as 95% accurate with tilt and azimuth within a 5% error margin [28] gives a buyer almost nothing to work with. Five percent of a 360-degree bearing is roughly 18 degrees, which may clear a structural-engineering tolerance and be useless for RF sector planning [42]. One accuracy number for a whole asset model is the wrong unit. The right unit is a per-parameter budget stated against the tolerance of the task that consumes it: mount height for structural loading, azimuth and downtilt for RF, cable and jumper routing for maintenance planning, ground-lead continuity for safety. Procurement should require accuracy per parameter, and the reference standard against which it was measured.

Acting when the model is known to be wrong. The DIRT dataset is the canonical case: 40,400 damages a year where the buried-asset model was wrong or unfindable, and a 38% chance across eight 811 centres of being unable to start on time because locate responses were incomplete [31,42]. A system that plans work from a registry inherits every error in it. The Charlot Lab position, applied here, is that each attribute in the plant model should carry a value, an uncertainty, a provenance and an age, and that every autonomous work order should carry an explicit predicate over those attributes that must hold for the plan to remain valid. When the predicate fails in the field, the correct output is abstention and escalation rather than a best guess. That makes refusal a first-class, measurable output. It also gives procurement a second metric alongside accuracy: the abstention rate. A vendor reporting 95% accuracy on the 60% of sites its system found tractable is not comparable to one reporting 90% across all of them, and today nothing in the public literature reports the denominator.

Contact. Everything deployed at scale in telecom is non-contact sensing. Drones photograph, twins measure, dispatch systems decide. Nothing fielded torques a bolt, dresses a connector or tests grounding continuity, and peer-reviewed work on robots that physically service towers remains at feasibility stage [41]. Built Robotics' trajectory is the useful signal: the best-funded autonomous trenching company found its volume business in solar pile driving rather than telecom conduit [37], because open repetitive greenfield geometry is what current autonomy handles and congested urban civils with surface-restoration constraints is not. NATE's framing of the near-term role, removing two or three climbs from a day of six to ten [13], is the honest ceiling for the next several years.

Currency. The single most valuable thing in the deployed record is not a robot. It is Verizon's policy of automatically capturing and digitising every new build and

modification from early 2025 [25], because a plant model decays continuously and a one-time capture programme produces an artefact that is wrong again within a maintenance cycle. Bentley's cost model explicitly excludes the cost of maintaining currency and of re-capture after each site modification [29], which means the published economics of asset twins omit the line item that determines whether the model stays useful. An operator evaluating this should price capture-on-change as the product and treat the initial sweep as the setup cost.

10. Where the value sits, ranked

Ranking by the size of the addressable pool, the quality of the evidence and what is technically and legally available now produces an ordering that inverts the emphasis of the vendor literature. The two applications with the strongest evidence require no embodied hardware at all. The application with the most marketing, structural tower inspection, ranks seventh.

The table below is the practical output of this review for a network operator. Sizes are drawn from the cited sources and, where marked, from the author's arithmetic on them. None of the readiness assessments should be read as a recommendation to buy or not to buy a specific product.

Rank	Opportunity	Size of the addressable pool	Best evidence	Readiness
1	Avoided and better-targeted dispatch	13% of on-site incidents need a second visit; 7.6 points of technician utilisation headroom, worth 9.2% more productive hours; $150–1,000 per roll	TSIA [15], Comcast [17,34], Vodafone [16,32], author's arithmetic [42]	Deployed at national scale; no embodied hardware required
2	Buildout labour productivity and method selection	Labour $12.23/ft underground; a 10% gain is $6,457 per mile; trenching to plowing is worth	FBA / Cartesian [1], Deutsche Telekom [36], author's	Planning available now; machine assistance in trial only [35,37]

		about 37.5% of method cost	arithmetic [42]	
3	Locate accuracy and damage prevention	About 40,400 damages a year from unmarked or mismarked facilities; 38% chance of a late start from incomplete locates	CGA DIRT [31], author's arithmetic [42]	A data and sensing problem; no autonomous system fielded
4	Plant-model currency (capture on change)	About $555M a year across the US tower estate if Bentley's per-tower model holds, roughly 1% of WIA cellular opex	Verizon [25], Bentley [29], author's arithmetic [42]	Process change available now; economics unaudited
5	Storm and mass-outage restoration	Not sized in any public source	Comcast [33]	Deployed, unmeasured
6	Field inventory and spares	No public figure found in this review	None located	Undocumented; the least-evidenced item in the fellow's own scope
7	Structural inspection automation	32,000–53,000 inspections a year; $21–36M a year at vendor prices, 0.04–0.07% of US cellular opex	TIA-222 [7], providers [8], author's arithmetic [42]	Deployed and small; gated on Part 108 for the crewless case [24]
8	Contact work on towers	Not sized	Academic literature [41]	Research stage

Table 4. Physical AI and logistics opportunities in telecom, ranked by evidenced value and current availability.

11. What is not known

The following are the open questions this review could not close. Each is stated with what would settle it, because several are answerable by an operator using data it already holds.

Field-operations labour as a distinct line. MTN Consulting reports $258.1bn of telecom labour across 4.36 million employees [19] and Omdia reports about $300bn [20], but neither isolates field crews from retail, care, IT and back office. Every claim in this subject about what automation is worth depends on a number nobody publishes. Any operator can produce it internally from payroll and cost-centre data, and doing so should precede any capital commitment.

The real cost of a truck roll. The $150-500 and $1,000 figures [14] have no traceable methodology and reach telecom through vendors selling avoidance. What is needed is a fully loaded figure per fault class, since a fibre splice, a small cell reboot and a residential install are different events with different costs, and the industry range averages across all of them.

Twin accuracy against physical truth. No study compares an asset digital twin against climb-verified measurement rather than against existing records [28]. Until one exists, twin accuracy claims are claims about agreement with the artefact whose reliability is the reason the twin was commissioned.

The guyed to self-supporting mix of the US tower estate. TIA-222 recommends three-year inspection for guyed masts and five-year for self-supporting structures [7], and the resulting national inspection volume varies by two-thirds depending on the mix, which is not published [42].

Whether telecom field labour is scarce. Continuum Capital projects a shortfall of 58,000 workers now and 119,000 over ten years [2]; WIA reports the wireless workforce falling 7.2% in a single year [4]; Wireless Estimator reports hiring at a twenty-year low [5]; BLS projects the technician occupation declining 3% [3]; CWA calls the shortage narrative misinformation [2]. These cannot all describe the same labour market, and no neutral source adjudicates.

The tower fatality rate. There is no BLS occupational code isolating tower technicians, so the published rates differ by 1.8x according to which denominator and which period is chosen [12]. The safety argument for removing people from towers cannot be quantified until a neutral series exists.

Cost of BVLOS compliance under Part 108. The final rule's detect-and-avoid and operator-certification requirements are not public [24], so the cost side of the crewless-inspection business case is unknown.

Positioning degradation near radiating structures. vHive's claim that EMI disrupts RTK enough to force manual flights [27] comes with no error figures and no degradation curve against transmit power or frequency. This is directly measurable with a survey-grade receiver and a spectrum analyser at a live site, and would be a genuinely useful piece of open work for the fellowship to produce.

Field inventory and spares economics. This is one of the four capabilities in the fellow's own project statement [43] and the review found no public figure for telecom field inventory carrying cost, stockout rate or its contribution to failed first visits. TSIA's 87% first-visit resolution [15] does not decompose into diagnosis failure, access failure and parts failure, and that decomposition is where the inventory case lives.

Whether any of it shows in the aggregates. Omdia notes that disclosures of quantifiable AI savings remain scarce [20], global telecom opex fell 0.2% in 2024 partly through divestment rather than efficiency [20], and Dell'Oro revised its own capital-intensity forecast by about a point within twelve months [18]. Any claimed automation effect is currently smaller than the disagreement between the houses measuring it.

12. Risk register

Regulatory schedule risk on BVLOS, high likelihood and high impact. The crewless-inspection business case depends on Part 108, which reached OIRA on 10 July 2026 with publication expected late 2026 or early 2027 and a six to twelve month transition, on no FAA commitment date [24]. A twelve-month slip pushes every associated payback out by the same amount. The compliance cost is unknown because the detect-and-avoid requirements are not public.

Labour-market reversal, moderate likelihood and high impact on the business case. Automation displaces expensive scarce labour more readily than cheap available labour. WIA reports a 7.2% workforce decline [4] and Wireless Estimator reports tower construction hiring at a twenty-year low with jobseeker registrations up more than 300% [5]. If that slack persists, the wage pressure that funds automation weakens even as the technology matures.

Vendor-sourced economics, high likelihood and moderate impact. The headline savings in this field are self-published or modelled. Bentley's $3,500 per tower is a model with three unreconciled percentages [29,42]; the $350-600 versus $800-1,500 inspection comparison is not scope-matched [8]; Comcast's tens of millions has no baseline, period or definition of an avoided roll [17,34]; Vodafone's numbers were published by its suppliers [16,32]. An operator that adopts vendor economics as its own business case inherits all of it.

Registry error compounding into autonomous action, moderate likelihood and high impact. A system that plans work from records containing a five-foot RAD centre error [28] or that digs on the basis of locate data that is wrong 20.5% of the time [31] will act confidently on bad inputs. This is the failure mode that turns an efficiency programme into an outage, and the mitigation is the uncertainty-and-abstention discipline in section 9 rather than better models.

Scope substitution in inspection, high likelihood and moderate impact. A drone pass and a climbing crew do different work. Torque checks, connector servicing, grounding continuity and remediation have no uncrewed equivalent [8,41]. Programmes that book the price difference as a saving without booking the residual crew visits will not realise the number.

Single-vendor dependency in the asset-model layer, moderate likelihood and moderate impact. vHive supplies the twin software behind at least two of the largest operator programmes, including Verizon and Deutsche Telekom [26], and Consilience Analytics has processed 35,000 twins into the Bentley stack [30]. The plant model is becoming a system of record, and concentration in the layer that holds it is a strategic exposure rather than a procurement detail.

Model-currency debt, high likelihood and moderate impact. Published twin economics exclude the cost of recapture and currency maintenance [29]. A programme funded on the initial sweep and not on capture-on-change produces an asset that decays into another unreliable record, which is where the industry started.

Capital tightening, high likelihood and moderate impact. Global capex was flat in 2025 and is forecast to fall 2% in 2026 [18], and MTN records the lowest twelve-month capex total since 2011 [19]. Verizon is guiding down while AT&T guides up [22]. Programmes that need multi-year capital commitments compete in a tightening envelope, which favours operating-expense-funded software over capitalised hardware fleets.

Measurement invisibility, moderate likelihood and moderate impact. If a programme's effect is smaller than the 16% disagreement between Omdia and MTN on telecom labour cost [19,20,42], the operator will not be able to prove it worked from aggregate financials and will need instrumented, matched-cohort measurement designed in from the start.

13. Scenarios to 2032

Scenario A, decision support carries the value. Part 108 publishes in late 2026 or early 2027 [24] but detect-and-avoid and certification costs keep per-site crewless economics marginal outside dense clusters of towers. Asset twins become standard practice through capture-on-change policies of the Verizon type [25]. Embodied field systems in the US stay marginal. Savings appear as dispatch avoidance and first-visit resolution, in the range the operator-attributed evidence already supports: 10% fewer dispatches [16], 28% fewer repeat visits [32], a few points of technician utilisation [15]. This is the base case on current evidence. Discriminating indicator: no US carrier publishes an audited before-and-after on robotic field work by the end of 2028.

Scenario B, the twin becomes the system of record. Operators and towercos mandate capture on every modification, and the twin displaces the RF data sheet and the structural drawing as the authority for what is on a structure. Value accrues through reduced rework, faster co-location approval and fewer failed site visits rather than through removed climbs. This is the scenario in which the fellow's project has the largest effect, because it makes the plant model a contractual object with accuracy obligations attached. Discriminating indicators: a carrier or towerco publishes per-parameter accuracy against a climb-verified baseline; lease and co-location contracts begin referencing the twin rather than the drawing; a vendor publishes an abstention rate alongside an accuracy figure.

Scenario C, robotics becomes a carrier revenue line rather than a cost line. The Chinese operator model generalises: China Telecom's Robotics-as-a-Service with AgiBot, China Unicom's hazard inspection deployments and State Grid's 8,500-robot programme [40] establish connectivity-plus-robotics as a sellable bundle, and Western carriers follow the pattern already visible in AT&T's Ghost Robotics work, which is a FirstNet product sold to public safety and defence rather than an internal tool [38]. Under this scenario the interesting question for an operator is not what robots can do to its own cost base but what its network can charge other people's robots. Discriminating indicators: a US or European carrier reports a robotics

services P&L line; FirstNet-style bundles are sold to utilities and industrial customers at volume.

A downside branch runs across all three. If the labour slack visible in WIA's 7.2% workforce decline [4] and Wireless Estimator's twenty-year hiring low [5] persists through 2028, the wage pressure funding automation weakens, and programmes that were justified on scarcity get repriced against cheap available crews. The indicators to watch are WIA's annual FTE figure and the postings-to-registrations ratio at [5], both of which are published yearly and neither of which requires access to proprietary research.

Across every scenario, one structural fact holds. WIA's own numbers put recurring operations at 5.4 times expansion capital in US wireless [4,42], and Dell'Oro has capex flat to slightly declining through 2030 [18]. The estate keeps growing while the capital to build it does not, which pushes the problem onto operations regardless of which scenario obtains.

14. Conclusions, and a measurement programme an operator can run

The economic opening the fellow's project is aimed at is real and the sources agree on its shape. Labour dominates buildout cost [1], recurring operations dominate wireless spend by more than five to one [4,42], costs are rising and 88% of respondents expect them to keep rising [1], and the plant model is demonstrably wrong at every scale from the national site count [4,6] down to a RAD centre height [28]. What the evidence does not support is the specific mechanism the market has been selling. Structural inspection automation, at national scale and at the vendors' own quoted prices, is worth a few tens of millions a year in the US [42]. The crewless flying it depends on ran at under 2% of flights in the FAA's own showcase programme [23].

The value that is actually documented sits in deciding correctly whether to send anyone. Comcast's dispatch prediction [17,34], Vodafone's remote assistance and technician decision support [16,32], and the 13% repeat-visit rate and 7.6 points of utilisation headroom in TSIA's benchmarks [15] are the evidenced pool. None of it requires a robot. All of it requires the same thing the embodied case requires, which is a plant model good enough to reason from and honest enough to say when it does not know.

That is where this review lands for a network operator, and it is a programme rather than a purchase. Six steps, in order, none of which requires a vendor commitment.

First, produce the field-operations labour number. Payroll and cost-centre data will give the fully loaded annual cost of field crews as a distinct line. No public source has it [19,20], and every subsequent decision is denominated in it.

Second, produce a truck-roll cost per fault class rather than a single average. The published $150-1,000 range [14] spans nearly sevenfold because it averages across events that have nothing in common. The decomposition of failed first visits into diagnosis failure, access failure and parts failure is the same exercise and answers the inventory question the public record cannot [15,43].

Third, run a climb-verified twin accuracy audit on a stratified sample of at least thirty structures, reporting agreement per parameter and adjudicating every disagreement by physical measurement. This is the study that does not exist anywhere in public [28], it is inexpensive relative to a twin programme, and it converts a vendor accuracy claim into a number that can be written into a contract.

Fourth, price capture-on-change as the product. Verizon's automatic capture of every new build and modification [25] is the mechanism that keeps a model current, and it is the line item the published economics omit [29]. Ask what the model costs to keep true for five years, not what the initial sweep costs.

Fifth, write two metrics into every procurement in this area: accuracy stated per parameter against a named reference standard, and abstention rate, meaning the fraction of sites or assets on which the system declined to produce a measurement, with accuracy reported on the remainder. A system that is right about what it knows and explicit about what it does not is the only kind that can be allowed to plan work autonomously against a registry that is wrong 20.5% of the time [31].

Sixth, hold BVLOS capital until Part 108 publishes and its detect-and-avoid requirements are readable [24]. The waiver route has cleared roughly 134 organisations in five years [23] and the compliance cost of the final rule is the largest unknown in the crewless-inspection case.

For the fellowship itself, the most valuable piece of open work visible from this evidence base is the positioning study. vHive asserts that electromagnetic interference around towers disrupts RTK badly enough to force manual flights [27] and publishes no measurement. A survey-grade receiver, a spectrum analyser and access to live and dark sites would produce a degradation curve against transmit

power and frequency that nobody has published and that every autonomous inspection programme needs. It is measurable, it is within the Charlot Lab's on-device state-estimation remit, and it would convert the central physical obstacle in this domain from a vendor claim into a public number.

References

1. [1] Fiber Broadband Association / Cartesian, "Fiber Deployment Cost Annual Report 2025", 3rd edition, published January 2026. Operator and contractor self-reports gathered by phone interview and online survey, September-October 2025, across projects in 38 states; each year's sample independent, so year-on-year comparison is not like-for-like; sample size not disclosed in the executive summary.
   REPORTED
2. [2] Continuum Capital, broadband workforce study commissioned by the Fiber Broadband Association and the Power & Communication Contractors Association, presented at Fiber Connect, July 2024. Forecast commissioned by trade bodies whose members benefit from workforce funding; predates the BEAD restructuring and the 2025 wireless construction contraction. The Communications Workers of America has publicly disputed US telecom workforce shortage claims.
   MODELLED
3. [3] US Bureau of Labor Statistics, Occupational Outlook Handbook, telecommunications equipment installers and repairers except line installers, 2024-34 projections. Projections are a model rather than a measurement; bls.gov blocked automated retrieval during compilation, so these figures come from search-index extracts of the OOH page rather than a direct read of the primary. Excludes line installers, who are projected to grow.
   MODELLED
4. [4] Wireless Infrastructure Association, "By the Numbers" 2025, research by iGR, released 18 March 2026. Association-commissioned and modelled by iGR from operator and contractor inputs; methodology and sample not public; figures are full-time equivalents rather than headcount.
   MODELLED (IGR ESTIMATE FOR WIA)
5. [5] Wireless Estimator, tower construction employment tracking against its own twenty-year record, article dated 5 February 2026. Proprietary single-source data from one industry job board, which is also the publisher's commercial base; postings volume is a noisy proxy for demand.
   REPORTED
6. [6] CTIA, Annual Wireless Industry Survey, 2025 edition, reporting year-end 2024. A survey of CTIA's own members with operator-supplied figures and no independent audit; CTIA's site definition differs from WIA's and the two counts are not reconcilable from published definitions.

REPORTED (MEMBER-SUPPLIED, TRADE ASSOCIATION)

7. [7] ANSI/TIA-222 (Telecommunications Industry Association), TR-14 committee FAQ on structural inspection intervals. TIA states explicitly that the interval is a recommendation rather than a mandate and that owners may lengthen or shorten it per structure.
   VERIFIED
8. [8] Drone tower inspection service providers (AeroDeploy, Drone Harmony, DroniTech, Voliro, ABJ Drones), 2025-2026 marketing material. Every figure traces to a company selling the service; no independent, operator-audited or peer-reviewed comparison was located; the drone and climbing-crew comparison is not scope-matched.
   SELF-PUBLISHED
9. [9] AT&T drone tower inspection programme, via Computerworld and RCR Wireless, 2016-2017. Nine to ten years old; no updated AT&T disclosure located; an AT&T official declined to state the cost. The 65,000-tower figure is the owned-and-operated estate at the time and predates subsequent tower sales.
   SELF-PUBLISHED
10. [10] American Tower and Crown Castle corporate communications on drone-derived asset digital twins, 2025. American Tower's industry-leader rating comes from its own customer survey with no published methodology, sample or comparison set; neither company discloses structures scanned, refresh cadence or model accuracy.
    SELF-PUBLISHED
11. [11] US Occupational Safety and Health Administration, Communication Towers topic page: 13 fatalities in 2013, 12 in 2014, 3 in 2015, 6 in 2016. Not updated past 2016 as of August 2026; counts only, with no published denominator, and no current official federal series exists for this work.
    VERIFIED
12. [12] Communications Workers of America survey (103 respondents) versus Wireless Estimator fatality database and cited safety experts, 2023. Both parties are interested; Wireless Estimator's database is proprietary and its own summary is internally inconsistent on its date range; there is no BLS occupational code isolating tower technicians, so no neutral rate exists.
    REPORTED
13. [13] National Association of Tower Erectors, via Commercial UAV News, on climb frequency and fatigue. An association's qualitative characterisation with no underlying measurement study cited and no data linking climb count to incident probability.
    REPORTED
14. [14] Technology & Services Industry Association truck-roll cost figure, as cited across field-service vendor literature (PTC, SightCall, Blitzz, VSight, Smarty), 2024-2026. No primary TSIA publication stating the $1,000 figure could be located; every

citation runs through a vendor selling truck-roll avoidance; the $150-500 band has no traceable methodology and mixes residential installs with enterprise service.
REPORTED

15. [15] Technology & Services Industry Association, field services benchmarking and "The State of Field Services 2025". Member-supplied benchmark data behind a membership gate, drawn from technology services generally rather than telecom field operations specifically; comparison companies are self-selected into the benchmark.
REPORTED

16. [16] Vodafone remote visual assistance results, published by the vendor TechSee, undated case study. Published by the supplier rather than by Vodafone, with no date, market, baseline or measurement period stated.
SELF-PUBLISHED

17. [17] Comcast dispatch-prediction machine learning system, via StreamTV Insider and Rethink Research. Comcast's own claim, with no baseline dispatch volume, no time period and no definition of what counts as an avoided roll.
SELF-PUBLISHED

18. [18] Dell'Oro Group, worldwide telecom capex forecast, 2026, covering about fifty service providers representing roughly 80% of global capex. Dell'Oro's own September 2025 release gave capex/revenue approaching 15% by 2029 and wireless intensity of 12-13%, so the house revised itself materially within a year; no absolute dollar total is published.
REPORTED

19. [19] MTN Consulting, Telco Market Tracker 2Q25, published December 2025, covering 140 operators. Varying disclosure quality across the sample; the labour figure aggregates all employees rather than field crews; figures are annualised trailing rather than calendar-year actuals.
REPORTED

20. [20] Omdia, Global Telecoms Opex Tracker. Omdia's roughly $300bn labour figure and MTN Consulting's $258.1bn describe the same thing on different samples and definitions; Omdia attributes part of the 2024 decline to divestments rather than efficiency and notes that disclosures of quantifiable AI savings remain scarce.
REPORTED

21. [21] Deloitte, 2026 Telecommunications Industry Outlook, citing IDC for global telecom revenue. Deloitte's capex intensity of about 18% is four points above Dell'Oro's projection, and its revenue base differs from Omdia's and MTN's; the outlook contains no field-operations or maintenance figures.
REPORTED

22. [22] AT&T, Verizon and T-Mobile capital investment guidance, 2025-2026. Company guidance rather than outturn; definitions of capital investment differ between the three carriers.
SELF-PUBLISHED

23. [23] US Department of Transportation Office of Inspector General, Report AV2025034, 30 June 2025, audit of FAA beyond-visual-line-of-sight efforts. Covers the BEYOND, PSP and IPA partnership programmes only; commercial waiver holders outside them are not counted. The OIG also found FAA's data validation process vulnerable to errors and mis-categorised at least one operator's observer status.
VERIFIED

24. [24] FAA Part 108 BVLOS rulemaking and OIRA docket status, as reported by trade press (Airdata, UAV HQ, Pilot Institute, DLA Piper), 2026. No FAA commitment date exists in any of the trade sources; OIRA timelines slip routinely; the final rule's detect-and-avoid and operator-certification requirements are not public.
REPORTED

25. [25] Verizon asset digital twin programme (David Wolff, VP Central Engineering; Candice DeMonteiro, Associate Director), via Inside Towers, 2025. Site count given as a range rather than a figure; no accuracy metric, cost per twin, or refresh cadence for the legacy estate; the 15-minute disaster-assessment claim is not independently verified.
SELF-PUBLISHED

26. [26] vHive Series B funding and twin production volume, via VentureBeat and company materials, 2025. Counts of twins generated conflate first captures, re-captures and non-telecom verticals; no breakdown by industry or customer is published.
SELF-PUBLISHED

27. [27] vHive announcement on electromagnetic interference, RTK degradation and autonomous capture on the DJI Matrice 4E, 25 November 2025, via DroneLife. No positioning-error figures, no degradation curve against transmit power or frequency, no mechanism description, and no third-party validation; the sole-provider claim is not independently checkable.
SELF-PUBLISHED

28. [28] Bentley Systems / Consilience Analytics validation whitepaper on OpenTower iQ, authored by Apurba Tribedi (Bentley), 2024. Vendor-sponsored; sample of four towers drawn from a group of fifty; the comparison baseline was best-existing data (RF data sheets, structural analyses and CAD drawings from under a year to several years old) rather than climb-verified ground truth; the paper does not state whether the 5% margin is of reading or of full scale.
SELF-PUBLISHED

29. [29] Bentley Systems operations and maintenance cost model within the same 2024 whitepaper: about $5,700 per tower manually against $2,200 with OpenTower iQ, plus 50% reporting time savings, 15% annual opex reduction and 42% lower ten-year TCO. Not a measurement; no operator has published a matched-cohort before-and-after; the model excludes twin currency maintenance, re-capture after site modification, and pilot and mobilisation cost.
MODELLED

30. [30] Consilience Analytics twin production volume, stated via the Bentley whitepaper, 2024: over 35,000 digital twins processed since 2019. Cumulative twins processed includes repeat captures of the same asset, so it is an upper bound on unique sites and cannot be read as estate coverage.
    SELF-PUBLISHED

31. [31] Common Ground Alliance, 2024 DIRT Report, released 28 August 2025, analysing 196,977 unique damage reports. A voluntary submission dataset that CGA itself treats as an undercount, with a reporting population that changes year to year; root cause is assigned by the reporting party, who has an interest in the assignment; telecom's share is inflated by telecom construction being a major cause of damage to other utilities; the five listed root causes sum to 67.9% against a summary characterisation of about 85%, a discrepancy this review could not resolve.
    REPORTED

32. [32] Vodafone and Celfocus, Field Technician Assist on Google Cloud, 2025 (FutureNet World 2025 awards finalist). Operator's own figure with no baseline period, market or fault mix disclosed and no independent audit; "repeated visits" is not defined in public materials.
    SELF-PUBLISHED

33. [33] Comcast alarm-clustering and tooling dispatch system for mass outages, trialled during the 2024 hurricane season and operational nationwide from September 2025. "Effectiveness" is not defined and the 50% claim has no stated denominator; no third-party verification and no per-storm breakdown.
    SELF-PUBLISHED

34. [34] Comcast dispatch-prediction accuracy and savings, via Rethink Research and an H2O.ai case study. Accuracy without a precision and recall split is close to meaningless for an imbalanced dispatch decision; the savings figure is unaudited and the timeframe is not stated.
    SELF-PUBLISHED

35. [35] Openreach / BT Group MicroBORE trenchless robotic drilling, developed with Synthotech and the University of Surrey, November 2025, via ISPreview, Telecompaper and IT Brief. No published bore rate, cost per drop, failure rate, trial volume, or comparison against manual moling; Openreach's £15bn programme to 25 million premises by December 2026 is being met without it.
    SELF-PUBLISHED (OPERATOR, VIA TRADE PRESS)

36. [36] Deutsche Telekom deployment cost per metre (about €85 traditional underground, €65 reduced-depth, €10 overground), via teltarif.de. Operator's own figures with no published methodology; German civils costs are dominated by local reinstatement and permitting rules that do not transfer to the US or UK.
    SELF-PUBLISHED

37. [37] Built Robotics Exosystem excavator retrofit, 2025, and the Blattner nationwide solar piling agreement reported by ENR, September 2025. The 20-30% saving is a

vendor-reported contractor figure with no named contractor or measured baseline; no telecom fibre-build deployment of the trenching product is publicly documented.
SELF-PUBLISHED

38. [38] AT&T and Ghost Robotics quadrupedal ground vehicles over FirstNet using AT&T Geocast, 2023 onward, via Light Reading and Urgent Communications. No fleet size, deployment count or telecom-facility use case disclosed; programme announcements date from 2023 with limited public updates since.
SELF-PUBLISHED

39. [39] NTT DATA data centre facility inspection robots, rolled out across Japanese sites from April 2023, via DataCenterDynamics and Digital Infrastructure Lab. 2023 vintage; no current fleet count, no measured labour saving, no evidence of extension to central offices or exchanges, and not independently verified.
SELF-PUBLISHED

40. [40] GSMA MWC Shanghai 2026 coverage of China Telecom's Robotics-as-a-Service with AgiBot and China Unicom's hazard-inspection deployments; State Grid Corporation of China robot programme announcements, 2025-2026. Announced programme figures from state-linked entities with no independent verification of units in service, uptime or task success rate; deployment targets of this kind routinely diverge from delivery.
SELF-PUBLISHED (OPERATOR, VIA EVENT COVERAGE)

41. [41] Shah and Dave, peer-reviewed review of robotic systems for communication tower inspection and maintenance, cited here only for the qualitative finding that such systems remain at feasibility stage. The full citation was truncated in the compiled evidence record for this review and must be completed before publication; no quantitative claim in this report rests on it.
REPORTED (LITERATURE, CITATION INCOMPLETE)

42. [42] Author's calculations. Every figure carrying this index is computed from the cited primaries rather than measured, and is presented as modelled. Derivations include: per-mile conversion of per-foot fiber costs at 5,280 ft/mi; labour share against total median cost and against labour plus materials; the residual between the sum of medians and the median total; per-mile value of the internal-versus-contractor labour delta and of a 10% labour productivity gain; the trenching-to-plowing saving implied by a 60% method spread; opex-to-capex ratio and per-site opex on WIA denominators; TIA-222 implied inspection volume across 158,500 towers at three- and five-year intervals; national value of drone structural inspection at quoted midpoints; extrapolation of Bentley's per-tower saving across the US tower estate; DIRT locate-error damage counts and root-cause sum; labour cost per employee from MTN aggregates; the Omdia-MTN labour gap; TSIA repeat-visit and utilisation headroom arithmetic; and decomposition of the tower fatality rate dispute into numerator and denominator contributions. Working was executed and checked programmatically on 5 August 2026.
MODELLED

43. [43] Ronnie Gomez, project statement, Industrial Research Fellowship, The Charlot Lab, Institute for Physical AI @ Bailey Military Institute, 2026. The fellow's own framing of scope and intended outcomes; used here to define the four capabilities the review tests and not as evidence for any quantitative claim.
SELF-PUBLISHED

Market and economic review, not investment advice. Figures carry the verification grade under each reference: **verified** (peer-reviewed or independently replicated), **reported** (a named source stated it, not independently confirmed), **self-published** (the organisation's own figure) and **modelled** (computed here, not measured). A modelled figure is never presented as a measurement. Corrections are welcome and will be recorded.