



AI IN TELECOM OPERATIONS

Five questions about AI in telecom operations

AI is the lever every operator names in its cost programme, and no annual report puts a currency figure on what it has saved. The three American carriers spent \$150.9 billion in 2025 running their networks and serving their customers, within one per cent of what they spent two years earlier. Read the five parts to find out:

1. the cost an AI saving is measured against and where it sits in the accounts
2. which savings operators have claimed, in which venue, and whether the accounts show them
3. ten candidates ranked by what it costs to get a number a finance director would accept
4. what 3GPP, O-RAN and the autonomy levels provide and what they leave out
5. where to put the kill gate on a pilot, and the two outages that show why

About this ebook

This ebook collects the five parts of *AI in Telecom Operations*, published by Blue Prysm at www.blueprysm.com/analysis/.

The website carries the current edition. Where a figure here differs from the website, the website is the later version.

EDITION

First edition, October 2026

DATA AS OF

PART 1 Forms 10-K for 2023 to 2025, Forms 10-Q for the first half of 2026, Vodafone's 20-F for the year to March 2026 and company sustainability data, read 7 Oct 2026

PART 2 filings, transcripts and reports from 2023 to the second-quarter 2026 results, read 7 Oct 2026

PART 3 operator, vendor and trade-press publications from 2020 to September 2026 and 3GPP TR 38.864 v18.1.0, read 7 Oct 2026

PART 4 3GPP specifications and release descriptions to TR 21.919 v19.0.0, O-RAN and operator publications to September 2026, read 7 Oct 2026

PART 5 rules, regulators' reports and trackers as read 7 Oct 2026

Method version 1.0

PUBLISHED BY

Blue Prysm LLC, Florida, United States · www.blueprysm.com · hello@blueprysm.com

© 2026 Blue Prysm. Free to use and adapt; attribution with a link to www.blueprysm.com is asked for, not required. Quotations from third-party documents remain the property of their authors and are cited where they appear.

Nothing in this ebook is legal, financial or investment advice.

Errors can be reported at www.blueprysm.com/contact/.

Contents

PART 01

The cost an AI saving is measured against

5

How large is the cost an AI saving is measured against, and where does it appear in the accounts?

Cost of services plus SG&A, which is where any operating AI saving has to land, came to **\$54.4B** at AT&T, **\$61.6B** at Verizon and **\$35.0B** at T-Mobile in 2025. The combined pool moved 0.8% in two years while revenue grew 5.1%. AT&T and Verizon employ **32,370** fewer people than they did two years ago, and their filings put that down to separation programmes and legacy shutdowns, with a severance charge beside each. Only Vodafone files an energy bill: **€0.9B** for 5,967 GWh.

PART 02

What operators say AI saves

16

Which AI savings have operators put a number on, in which venue, and can the accounts show them?

None of the Forms 10-K and 20-F or annual reports of the eleven operators in this series puts a currency figure on what AI saves. The programme totals treat AI as one lever. AT&T targets **\$4B** a year by the end of 2028 and Verizon **\$5B** of 2026 opex, “with a substantial portion realized by headcount reductions”. T-Mobile expects nearly \$3B from “digitalization and AI” by 2027 and BT £3.7B of savings by FY30. The figures that name AI were said on calls, such as Verizon’s **\$200M** of energy with no period and no base, or appear as a share, such as the **34%** of enquiries Telstra’s assistant handled.

PART 03

Ten AI candidates ranked by the cost of finding out

27

Which AI candidates can produce a trustworthy number cheapest, and what has been published for each?

Energy saving ranks first of ten, and it is the only candidate for which an operator has published a national figure. BT’s cell sleep on **19,500** sites saves 4.5 GWh a year, which works out at about **0.63 kWh** per site per day against a headline of up to 2 kWh. 3GPP’s own simulations put one technique’s saving anywhere from **0.9% to 84.8%**, depending on load. Care has results by journey and no all-contact containment rate with a definition behind it. Field dispatch and procurement have no published baseline at all, and they rank last.

PART 04

What the standards add

38

What do 3GPP, O-RAN and the autonomy levels give an operator to build an AI programme on?

The 5G core has had an analytics function, NWDAF, since Release 15, but a standard way to monitor whether a model is still accurate arrived only in March 2023, in TS 23.288 v18.1.0. Release 17 limited model sharing to **single-vendor** environments. The radio study found **no consensus** for AI-based CSI compression, and Release 19 standardised the one-sided models a single vendor can build, implementable from December 2025. One third-party rApp runs in a live network, AT&T’s since July 2025, and no result from it has been published. China Mobile reports L4 in **21 high-value scenarios**, against its 2023 target of L4 by 2025 for “almost all the services”.

PART 05

Where the kill gate goes

Which rules and which failures decide where an AI pilot must stop, and what should it measure while it runs?


At Rogers in 2022 “the risk assessment algorithm downgraded the risk level” to Low on the change that cut off over **12 million** customers. After AT&T’s outage of 22 February 2024, which blocked more than **92 million** calls, the FCC wrote that loading such changes “should not be possible”. AI voices need prior express consent under FCC 24-17, EU chatbots have had to say what they are since 2 August 2026, and Regulation 2026/1744 moves the high-risk duties for critical infrastructure to **2 December 2027**. The first gate goes where no automated tool can lower a change’s risk class.

Work with Blue Prysm

The cost an AI saving is measured against

How large is the cost an AI saving is measured against, and where does it appear in the accounts?

Any claim about what AI saves a carrier needs a base to be measured against. AT&T, Verizon and T-Mobile spent \$150.9 billion in 2025 on running their networks and on selling to and serving their customers, within 1% of what they spent two years earlier. An AI saving has to show up as a fall in one of those lines. So far the one line that has clearly moved is headcount, and the carriers credit separation programmes and the shutdown of old networks rather than any algorithm.



IN BRIEF

- Cost of services plus SG&A, the two lines an operating AI can lower, came to **\$150.9 billion** across AT&T, Verizon and T-Mobile in 2025 (Forms 10-K). That is 0.8% more than in 2023, over two years in which their revenue grew 5.1%.
- What has moved is headcount. AT&T ended 2025 with 133,030 employees, **16,870** fewer than at the start of 2024, and Verizon with 89,900 full-time equivalents, **15,500** fewer than two years earlier. Each filing credits a separation programme or a legacy shutdown, and each records a severance charge beside it.
- No carrier files what customer care costs, how many contacts it handles or what it pays for electricity. Vodafone alone publishes an energy spend, **€0.9 billion** for 5,967 GWh in FY26, which makes a kilowatt-hour the one AI saving a reader can check against a meter.

The \$150.9 billion operating cost pool has hardly moved in two years

An operating AI can only save money in two lines of a carrier's income statement. One is the cost of running the network and serving customers, which Verizon and T-Mobile call cost of services. The other is the selling, general and administrative expense around it. For AT&T, Verizon and T-Mobile the two lines together came to \$150.9 billion in 2025 (Exhibit 1).


The captions differ from one carrier to the next. Verizon and T-Mobile report cost of services; AT&T calls the same line other cost of revenues. All three leave depreciation and amortisation out of it. Verizon's own definition of cost of services runs from "salaries and wages" through "network access and transport costs" to "computer systems support". Its SG&A takes in sales commissions, advertising, bad debt and "call center and information technology costs". Between them, those two lines hold every cost an operating AI could plausibly touch.

Two costs are left out. Equipment cost rises and falls with device sales and says nothing about how the operation is run. Depreciation does not change in the year an operating AI is switched on.

EXHIBIT 1

Cost of services and SG&A came to \$150.9 billion across the three carriers in 2025, and no carrier's pool moved more than 6.1% in two years

■ Cost of services ■ SG&A ■ 2023 total



\$ billions, as filed. Equipment cost, depreciation and amortisation are left out.

Source: AT&T, Verizon and T-Mobile Forms 10-K for 2023 and 2025, consolidated statements of income (Filed); the pool is the sum of the two lines (Calculated). Note: AT&T's cost of services is the line it captions other cost of revenues.

The three carriers went in different directions. AT&T's pool fell \$1.5 billion to \$54.4 billion while its revenue rose 2.6%; its other cost of revenues came down 5.8% and its SG&A rose 0.2%. Verizon's pool went the other way, up 1.3% to \$61.6 billion on revenue that grew 3.1%. Excluding Verizon's own special items, its cost of services fell 1.5% and its SG&A rose 0.8%. T-Mobile's pool rose most, by 6.1% to \$35.0 billion, but its revenue rose 12.4% and its 2025 figures take in UScellular from August of that year (Exhibit 2).

EXHIBIT 2

The operating cost pool was 43.3% of revenue at AT&T, 44.6% at Verizon and 39.6% at T-Mobile in 2025

\$ MILLIONS, AS FILED	AT&T	VERIZON	T - MOBILE
Revenue, 2025	125,648	138,191	88,309
Cost of services, 2025	25,424	27,789	11,497
SG&A, 2025	28,942	33,818	23,470
Operating cost pool, 2025	54,366	61,607	34,967
Operating cost pool, 2023	55,861	60,845	32,966
Change in the pool, 2023 to 2025	-2.7%	+1.3%	+6.1%
Change in revenue, 2023 to 2025	+2.6%	+3.1%	+12.4%
Pool as a share of revenue, 2025	43.3%	44.6%	39.6%
Employees at year end, 2025	133,030	89,900 FTE	about 75,000
Pool per employee, 2025	\$409k	\$685k	\$466k

AT&T's cost of services is the line it captions "Other cost of revenues (exclusive of depreciation and amortization)". The pool leaves out equipment cost, depreciation and amortisation and impairments. The carriers count employees on different bases, so the last line is indicative only.

Source: AT&T, Verizon and T-Mobile Forms 10-K for 2023 and 2025, statements of income and human-capital sections (Filed); the pool, the changes and the ratios are calculated from the filed lines (Calculated).

A narrower pool gives the same answer. No operating AI runs advertising or decides who defaults. Without those two lines, and without Verizon's special items, the pool was \$49.0 billion at AT&T, \$53.0 billion at Verizon and \$29.9 billion at T-Mobile in 2025. Two years earlier the same lines came to \$51.3 billion, \$53.3 billion and \$29.6 billion.

0.8%

How much the three carriers' combined operating cost pool moved between 2023 (\$149.7 billion) and 2025 (\$150.9 billion). Their revenue grew 5.1% over the same two years.

Whatever AI has saved so far is buried inside a cost base that held level in dollars and fell as a share of revenue. The only way a claimed saving can be checked is as a movement in one of these lines, which is the test the rest of the series applies.


AT&T and Verizon have 32,370 fewer people, and neither cites AI

The part of the pool that has moved most is people. AT&T ended 2025 with 133,030 employees, 16,870 fewer than at 31 January 2024. Verizon ended it with 89,900 full-time equivalents, 15,500 fewer than two years earlier. T-Mobile went the other way, from about 67,000 to about 75,000, because it was buying businesses (Exhibit 3).

EXHIBIT 3

AT&T and Verizon together employed 32,370 fewer people at the end of 2025 than two years earlier, while T-Mobile grew as it bought businesses

■ End 2023 ■ End 2025



Source: AT&T, Verizon and T-Mobile Forms 10-K for 2023 and 2025, human-capital sections (Filed). Note: AT&T's 2023 figure is at 31 January 2024. The three carriers count on different bases, so the rows compare each carrier with itself and not with the others.

Each carrier gives a reason, and in every case it is a programme rather than a technology.

- **Verizon.** By the end of March 2025 about 4,800 people had left Verizon under a voluntary separation programme. In the fourth quarter it announced a cut of more than 13,000 positions, and “Over 80% of the affected employees exited in December 2025”. The severance for the year came to \$1.5 billion before tax. The [8-K of 20 November 2025](#) that announced the cut does not mention AI. It names “a significant reduction in outsourced, contracted and other outside labor expense”.
- **AT&T.** AT&T's 10-K for 2025 explains lower Business Wireline costs as “primarily driven by lower personnel and customer support costs associated with ongoing transformation initiatives”. Its 10-Q for the second quarter of 2026 cites “lower personnel and other costs resulting from the decommissioning of our legacy network”. Neither filing mentions AI.
- **T-Mobile.** T-Mobile cut about 5,000 positions in August 2023, “primarily in corporate and back-office functions”. In late 2025 it began what it calls a workforce transformation, taking out layers of management and duplicated roles, and booked a \$390 million charge for it.

VERIZON: THE SAVING REACHES THE ACCOUNTS A YEAR AFTER THE CHARGE

Verizon's 2025 accounts show that a saving and its charge arrive a year apart. They record the severance for the 13,000 positions cut in 2025 alongside the saving from the people who left the year before.

That earlier saving is itemised, which is rare. Personnel cost fell \$222 million in cost of services “due to prior year workforce reductions”, and a further \$115 million in SG&A.

The saving from the December 2025 exits will not show until the 10-K for 2026, due in February 2027. [Part 2](#) uses that filing as the first test of the carriers' 2026 targets.

The pattern holds outside the United States (Exhibit 4). Deutsche Telekom's German headcount fell 4.9% in 2025, “due in particular to efficiency enhancement measures”. BT's labour resource, counting contractors, “fell 7% year-on-year to below 108,000”. Telstra cut 1,219 direct roles in FY26 and saw its labour expenses fall 4.6% to A\$3,807 million. Vodafone is the exception in the table, with more employees at the end of FY26 than a year before.

EXHIBIT 4

Deutsche Telekom, BT and Telstra report fewer employees in their latest year; Vodafone reports more

OPERATOR	MEASURE	EARLIER	LATEST	WHAT THE REPORT SAYS
Deutsche Telekom	Average employees, group, including T-Mobile US	204,856 (2023)	199,223 (2025)	Germany down 4.9% in 2025, “due in particular to efficiency enhancement measures” and staff restructuring
BT Group	Employees at year end	94,135 (FY24)	79,390 (FY26)	Labour resource, contractors included, “fell 7% year-on-year to below 108,000”
Vodafone	Employees at 31 March	87,007 (FY25)	91,128 (FY26)	Average of 92,988 in FY26, of whom 56,922 in customer care and administration
Telstra	Direct roles	30,553 (FY25)	29,334 (FY26)	Labour expenses down 4.6% to A\$3,807 million

BT's and Vodafone's years end on 31 March, Telstra's on 30 June. Telstra's FY25 figure is FY26 plus the 1,219 roles it reports cutting.

Source: Deutsche Telekom Annual Report 2025, note 25; Vodafone Form 20-F FY2026, note 24; Telstra Financial Results and Annual Report FY26; BT Annual Report 2026 for the labour-resource line (Filed). BT's employee counts are from its Responsible Business Addendum 2026 (Reported).

None of the Forms 10-K, 10-Q or 20-F behind this series attributes a fall in headcount to AI. Where a filing explains a fall, it names a separation programme, a restructuring or the shutdown of a legacy network, and the American ones put a severance charge beside it. AI may have played a part, but no filing says so.

No carrier files what customer care costs

Customer care appears in every carrier's statement about AI and in none of their accounts as a line. Verizon is the only one that says where the cost goes: "Aggregate customer service costs, which include billing and service provisioning, are allocated between Cost of services and Selling, general and administrative expense". The cost is in the pool, split two ways, and invisible from outside the company.

Operators disclose proxies instead, and they vary widely (Exhibit 5). Vodafone's Form 20-F is the one filing that counts care staff: 56,922 of an average 92,988 employees in FY26 worked in customer care and administration, which works out at 61.2% of everyone it employed. Telstra's annual report is the only one to give an AI share of contacts: its AI Assistant has handled 34% of enquiries since March 2026. T-Mobile reports "more than 50% reduction in calls to care since 2021" without saying how many calls that was. The nearest thing to a base is Light Reading's 2022 figure of 3 million calls a week.

EXHIBIT 5

Operators publish shares, counts and minutes for customer care, and none files its care cost or a current contact volume

OPERATOR	MEASURE	VALUE	WHERE	CLASS
Vodafone	Average employees in customer care and administration	56,922 (FY26), 61.2% of all employees	Form 20-F, note 24	Filed
T-Mobile	Calls to care	"more than 50% reduction" since 2021	8-K exhibit, 11 Feb 2026	Filed
T-Mobile	Calls to care, 2022	3 million a week, about 156 million a year	Light Reading , 19 Sep 2024	Reported
Telstra	Enquiries handled by its AI Assistant since March 2026	34%	Annual report, FY26	Filed
Deutsche Telekom	Front-end staff in Germany with direct customer contact	almost 27,000	Annual Report 2025	Filed
Deutsche Telekom	First-call resolution	76.0% (2025), 74.1% (2024)	Annual Report 2025	Filed
BT	Time customers spend on the phone to BT	down "almost three minutes"	Annual Report 2026	Filed
Verizon	Service team	28,000 people	Reuters , 9 Apr 2025	Reported

Vodafone's share and T-Mobile's annual figure are calculated from the published numbers.

Source: the filings and reports named in each row, classed by venue (Compiled).


A care saving stated as a percentage therefore has no filed base, because no operator here files either its contact volume or its care cost. The care figures in [part 2](#) are shares rather than sums, and will stay that way until an operator files a base.

Energy is the only cost with a physical unit

Energy differs from every other cost in the pool because it is counted in watt-hours before anyone puts a price on it, which makes it the easiest saving in the business to check. Every operator in this set publishes its watt-hours. Only Vodafone says what they cost (Exhibit 6).

EXHIBIT 6

Six operators publish their energy in watt-hours, and Vodafone alone files the bill: €0.9 billion for 5,967 GWh



Source: AT&T corporate responsibility data for 2025; Verizon EY review report for 2025, converted from 41,076,950 GJ at 3,600 GJ to the GWh; T-Mobile 2025 data tables; Deutsche Telekom CR report 2025 (Reported). Vodafone Form 20-F for FY2026, non-financial KPI table; BT Annual Report 2026, page 71 (Filed). Note: Deutsche Telekom's total includes T-Mobile US, so the two rows overlap. 1 TWh is 1,000 GWh.

Vodafone's figure is useful because it prices the unit. Its €0.9 billion for 5,967 GWh works out at about €150 a megawatt-hour, an average across every market it operates in, Africa included, on a spend rounded to the nearest tenth of a billion. The nearest thing to an American price comes from AT&T's own energy projects. AT&T says they saved \$81.1 million a year for 709 million kWh conserved in 2025, which is about 11.4 cents a kilowatt-hour. The same calculation gives between 8.7 and 11.4 cents in each year since 2022.

AT&T'S ENERGY PROJECTS, 2022 TO 2025

AT&T reports two numbers for its energy projects each year: the annualised saving in dollars and the kilowatt-hours conserved. For 2025 the projects saved \$81.1 million for 709 million kWh, about 11.4 cents a kilowatt-hour.

Set the 709 million kWh against AT&T's total consumption of 15.0 TWh and the projects conserved about 4.7% of what it used in 2025.

AT&T also describes using "AT&T-patented machine learning (ML) algorithms and automation" to let cell-site capacity sleep when load is low. It attaches no saving to that work, so its share of the 709 million kWh is unknown.

The totals themselves have moved little, which matters for anyone claiming a large energy saving. AT&T used about 6% less energy in 2025 than in 2022. T-Mobile, on a series recast for its acquisitions, used 14.9% less, and the energy it uses per petabyte carried fell from 263 MWh to 86 MWh. Verizon's assured total for 2025 is 0.2% below its 2022 figure on the same SASB definition. Vodafone went up, from 5,701 GWh in

FY24, restated to include Three UK, to 5,967 GWh in FY26. BT came down, from 2,989 GWh in FY25 to 2,678 GWh.

Price appears in only one place in an American 10-K. Verizon's Consumer segment reported a \$75 million rise in building and facility costs in 2025, "primarily due to higher utility rates". None of the three carriers publishes its energy bill, so a dollar saving on energy is a claim about a number the reader cannot see. A kilowatt-hour saving can be checked against the published total, which is why [part 3](#) ranks energy first.

Terms used

- **Cost of services.** The cost a carrier attributes directly to providing a service: wages, network access and transport, contracted services and systems support. AT&T's equivalent line is other cost of revenues.
- **SG&A.** Selling, general and administrative expense: sales commissions, advertising, bad debt, administration and, in Verizon's definition, call centres and information technology.
- **Operating cost pool.** The term used in this series for the two lines added together, which is the base any operating saving is measured against.
- **Special items.** Charges a company separates from its underlying results, such as severance and asset write-downs. Verizon tabulates them in its 10-K.
- **Full-time equivalent.** Headcount expressed as full-time positions, so two half-time employees count as one.

Implications

CARRIER STRATEGIST

State an AI target as a change in a named line of Exhibit 2, by year. A saving that cannot be found in cost of services or SG&A cannot be audited, and will not be believed. Your peers' 2026 claims will be tested against these figures for the first time in the Forms 10-K for 2026, and so will yours.

INVESTOR

Read any AI saving as a share of the carrier's own pool. The combined pool is \$150.9 billion and it moved 0.8% in two years, so a claim in the billions is a claim about a line you can watch. Take the headcount falls as the carriers describe them, as separation programmes and legacy shutdowns with a severance charge beside each, until a filing says otherwise.

VENDOR

Lead with energy, because it is the one cost a customer can verify with a meter. That makes it the easiest AI case to close and the most exposed to a measured result. Expect to be asked for site kilowatt-hours against a control, and have them ready.

Method and limits

HOW THIS WAS BUILT

Every financial figure is taken from the consolidated statement of income and the notes of each carrier's Form 10-K for 2023, 2024 and 2025, and checked against the XBRL company facts at data.sec.gov. The operating cost pool is cost of services plus selling, general and administrative expense, both before depreciation and amortisation. AT&T captions the first of those lines other cost of revenues. Headcount is the figure each 10-K states in its human-capital section, on whatever basis it uses.

Energy comes from each company's sustainability data, from Vodafone's 20-F and from BT's annual report. Verizon's gigajoules are converted at 3,600 GJ to the GWh. AT&T's price per kilowatt-hour is calculated as its annualised savings from energy projects divided by the energy they conserved.

WHAT IT DOES NOT SHOW

The carriers do not count people the same way. AT&T reports persons, Verizon full-time equivalents and T-Mobile full- and part-time heads, and AT&T's 2023 figure is taken at 31 January 2024. Acquisitions move the 2025 and 2026 bases as well: T-Mobile consolidated UScellular's wireless business from August 2025, Verizon closed Frontier on 20 January 2026 and AT&T bought Lumen's mass-markets fibre business in 2026. No carrier files its customer-care cost, its network-access cost or its energy cost in dollars, so the pool cannot be divided by function. The energy figures are sustainability data rather than audited accounts. Verizon's totals for 2023 and 2024 could not be found, and its 2022 figure excludes TracFone, which its 2025 operational-control boundary takes in.

Data as of: Forms 10-K for 2023 to 2025, Forms 10-Q for the first half of 2026, Vodafone's 20-F for the year to March 2026 and company sustainability data, read 7 Oct 2026 · Method version 1.0.

Sources. Forms 10-K of AT&T for 2025 and 2023; Verizon for 2025 and 2023; T-Mobile for 2025 and 2023, with each carrier's 10-K for 2024 and SEC XBRL company facts; AT&T's 10-Q for the second quarter of 2026; Verizon's 8-K of 20 November 2025; T-Mobile's results release of 11 February 2026; AT&T's corporate responsibility data and energy page; Verizon's EY review report for 2025 and 2023 ESG report; T-Mobile's 2025 data tables; Deutsche Telekom's Annual Report 2025, note 25, strategy section and CR report on energy; Vodafone's Form 20-F for FY2026; BT's Annual Report 2026 and Responsible Business Addendum 2026; Telstra's FY26 annual report; Light Reading and Reuters as named in Exhibit 5. Read 7 Oct 2026.

What operators say AI saves

Which AI savings have operators put a number on, in which venue, and can the accounts show them?

Eleven large operators have put a number next to AI since 2023, and none of the annual reports, 10-Ks or 20-Fs behind this series says what AI saves in currency. The large figures are programme totals in which AI stands beside headcount, property and contracts as one lever of several. The figures that credit AI by name were spoken on earnings calls, and the only one that made it into an annual report is Telstra's share of enquiries.

BLUE PRYSM · ANALYSIS

56 mentions of AI · 0 dollar figures

IN BRIEF

- None of the Forms 10-K and 20-F or annual reports of eleven operators puts a currency figure on savings from AI. The three American carriers' 10-Ks mentioned AI **56** times for 2025, up from 11 for 2023, and none of those mentions has a dollar amount next to it.
- The large numbers are programme totals that bundle AI with headcount, property and contracts. AT&T targets **\$4 billion** a year by the end of 2028, and Verizon **\$5 billion** of 2026 opex, "with a substantial portion realized by headcount reductions". T-Mobile expects nearly \$3 billion by 2027 from "digitalization and AI", and BT £3.7 billion of savings by FY30.
- The figures that credit AI by name were spoken on earnings calls. Verizon's "over \$200 million" of energy savings comes with no period and no base, and T-Mobile's 60% containment rests on one transcript. Telstra's **34%** of enquiries is the one AI share anywhere in an annual report.
- Verizon's \$5 billion is 8.1% of its 2025 operating cost pool and AT&T's \$4 billion is 7.4% of its own. The first filed test comes with the Forms 10-K for 2026, due in February 2027. Until then, an investor should sort each figure by where it was said before modelling it.

The venue of an AI claim decides how hard it can be held

An operator talking about AI savings is making one of three kinds of claim. The first is a programme total in which AI is one named lever. The second is an AI-specific operating figure, such as the share of contacts a bot handled. The third is a figure for AI revenue. Any of the three can be made in a filing, on an earnings call or to the press. In our view the venue says more about how hard the number can be held than the number itself does.

A filing is the company's document of record, and anything in it can be traced to the accounts. A call is on the record too, but it reaches most readers through third-party transcripts, and what is said on it is a forecast as often as a result. A release or an interview is attributed but unchecked. Sorted that way, the statements of eleven operators are lopsided. The large numbers are programme totals, and the AI-specific ones are small, spoken or both (Exhibit 1).

EXHIBIT 1

Across eleven operators, the largest AI figures are programme totals in which AI is one lever, and no annual report or SEC filing states an AI saving in currency

OPERATOR	LARGEST FIGURE	SPECIFIC TO AI?	VENUE AND DATE	CLASS
AT&T	\$4B a year of cost savings by end-2028, after over \$1B in 2025	No: one lever of several	Q4 2025 call, 28 Jan 2026, two transcripts	Filed
Verizon	\$5B of 2026 opex savings	No: "a substantial portion" from headcount	Q4 2025 call, 30 Jan 2026, two transcripts	Filed
Verizon	Over \$200M of energy savings; 85% of issues resolved autonomously	Yes	Q1 2026 call, 27 Apr 2026, with press	Filed
T-Mobile	Nearly \$3B of Core Adjusted EBITDA in 2027 against 2025	Partly: "digitalization and AI"	8-K exhibit, 11 Feb 2026	Filed
T-Mobile	About 60% chatbot containment	Yes	Q1 2026 call, 28 Apr 2026, one transcript	Reported
BT	£3.7B of cost savings by FY30	No	Annual Report 2026	Filed
BT	About 10,000 roles to AI and digitisation by 2030	Yes, in roles	Results-day remarks, 18 May 2023	Reported
Telstra	AI Assistant handled 34% of enquiries	Yes: a share	Annual report, FY26	Filed
Telefónica	€2.3B of gross efficiencies by 2028, €3.0B by 2030	No: one lever of several	6-K, 4 Nov 2025	Filed
Orange	More than €600M of "value" from AI by 2028	Yes; cost and revenue not split	Company web page	Reported
Vodafone	Over €100M of net benefit in FY26	Yes	CTO to TM Forum Inform, 22 Sep 2026	Reported
Deutsche Telekom	None; a cooling test showed potential "of up to 33 %"	n/a	CR report 2025	Reported
China Mobile	AI services revenue RMB90.8B in 2025	Revenue, broadly defined	Annual Report 2025	Filed


Source: the documents in the sources line, each row classed by its venue (Compiled). Note: a call figure is Filed where two transcripts, or a transcript and a same-week report, agree; a figure in one third-party transcript is Reported.

The 10-Ks tell the same story from the other side. AT&T, Verizon and T-Mobile mentioned AI 11 times between them in their 10-Ks for 2023 and 56 times in the ones for 2025, five times as often. Not one of the nine documents attaches a dollar amount to any of those mentions (Exhibit 2). AT&T's 10-K for 2025 repeats its 10-K for 2024 almost word for word. It expects "cost savings through AI-driven efficiencies in network design and operations, software development, sales, marketing, customer support services and general and administrative costs". Verizon's adds a caution. There "can be no assurance that the usage of AI will meaningfully enhance our products or services or be beneficial to our business, including our efficiency or profitability".

EXHIBIT 2

The three carriers' 10-Ks mentioned AI 56 times for 2025, five times as often as for 2023, and no mention comes with a dollar figure

■ FY2023 ■ FY2024 ■ FY2025



With a dollar figure: 0

With a dollar figure: 0

With a dollar figure: 0

Mentions of "AI" or "artificial intelligence"; a defined pair counts once.

Source: AT&T, Verizon and T-Mobile Forms 10-K for fiscal 2023 to 2025, main document, all sections, exhibits excluded (Filed); counts by Blue Prism of each whole-word "AI", hyphenated forms included, and each "artificial intelligence" in any case, with "artificial intelligence (AI)" counted once (Calculated). Note: no sentence in the nine filings states an amount beside AI.

The large numbers are programme totals with AI as one lever among several

Every large number an operator has published is a programme total. Each lists AI beside other levers, and none says how much of the total AI is supposed to deliver.

- **AT&T.** AT&T first tied AI to cost in its 10-K for 2023, where it expected “cost savings through AI-driven efficiencies in our network design, software development and customer support services”. By the middle of that year it had already met a \$6 billion run-rate savings target. Its investor-day release of [3 December 2024](#) then set “\$3 billion+ in run-rate cost savings by the end of 2027” and did not mention AI. On 28 January 2026 the chief financial officer said AT&T had “achieved over \$1 billion of cost savings in 2025” and expected “an additional \$4 billion annual cost savings by the end of 2028”. In April, in the one third-party transcript, he listed the levers: workforce optimisation and vendor rationalisation, AI enablement, digitalisation and lower legacy support costs. He gave no split between them.
- **Verizon.** Verizon’s target is the largest of the three, and it is the plainest about what is in it. On 30 January 2026 its chief executive described “an in-year war chest of \$5 billion in OpEx savings, with a substantial portion realized by headcount reductions”. The rest of his list was “marketing efficiencies, real estate rationalization, contract renegotiations and more”. The chief financial officer added “reducing call volumes, or on the IT side, in terms of rationalizing platforms and including AI enablement”, and gave no AI share. The [8-K](#) on the 13,000 job cuts does not mention AI.
- **T-Mobile.** T-Mobile is the one carrier with AI in the headline of a furnished filing. Its results release of 11 February 2026 promises “a nearly \$3 billion incremental contribution to Core Adjusted EBITDA expected from digitalization and AI, by the end of 2027, relative to 2025”. In the guidance section of the same release the credit goes to “significant 2025 investments in greenfield network expansion, AI, and digitalization”, so the attribution shifts inside a single document. On the call, in the one third-party transcript, the chief financial officer put the savings at \$1.3 billion in 2026 and \$2.7 billion in 2027, both against 2025.
- **BT.** BT’s filed number is its cost programme: £1.5 billion realised so far and an extension to FY30 “to enable the delivery of £3.7bn of cost savings”. Beside it is the line that “AI will continue to change how we work and the types of roles we need”. BT’s AI figure is older, and it was never filed. On 18 May 2023 AP, The Guardian and RCR Wireless reported that about 10,000 of the up to 55,000 roles to go by 2030 would go “through digitization, automation and the use of AI”. The Annual Report 2026 repeats neither number.
- **Telefónica.** Telefónica’s capital markets deck of November 2025 lists an “Autonomous Network Journey” and a “B2C call centers AI virtual assistant” among the levers behind €2.3 billion of gross efficiencies by 2028. There is no split by lever. Its exit plan for about 5,500 employees does not name AI.

FILED


AT&T call of 28 January 2026 and Verizon call of 30 January 2026, each in two transcripts (MarketBeat and The Motley Fool) · AT&T and Verizon 10-Ks and 8-Ks · T-Mobile 8-K exhibit of 11 February 2026 · BT Annual Report 2026 · Telefónica 6-Ks of 4 November and 22 December 2025

AT&T's list of levers, call of 22 April 2026, and T-Mobile's \$1.3 billion and \$2.7 billion, call of 11 February 2026, each in one third-party transcript (MarketBeat) · BT's 10,000 roles, AP, The Guardian and RCR Wireless, 18 May 2023

Set against the cost pool in [part 1](#), each American target comes into proportion. Verizon's \$5 billion is 8.1% of its 2025 pool of \$61.6 billion. AT&T's \$4 billion is 7.4% of \$54.4 billion. T-Mobile's \$1.3 billion for 2026 is 3.7% of \$35.0 billion, and its nearly \$3 billion for 2027 would be 8.6% (Exhibit 3).

EXHIBIT 3

Each American target is 4% to 9% of the carrier's 2025 operating cost pool, and none states an AI share



Source: targets as stated in Exhibit 1, divided by the 2025 operating cost pool in part 1 (Calculated). Note: a gross saving shows as a fall in the pool only where nothing else grows. Frontier, closed on 20 January 2026, adds its cost base to Verizon's, and UScellular is in T-Mobile's from August 2025. T-Mobile's \$3 billion is an EBITDA contribution and may include revenue.

Every programme total the eleven operators have published bundles AI with something else, whether headcount, property, contracts, legacy shutdowns or digital channels. None says how much of it is AI. A reader can hold AT&T to \$4 billion and Verizon to \$5 billion, but neither can be held to an AI figure because neither has given one.

Only one figure that credits AI has reached an annual report

Only a handful of figures credit a result to AI by name, and all but one were said on an earnings call rather than written into a report.

\$200M

VERIZON: “OVER \$200 MILLION” OF ENERGY SAVINGS FROM AI, Q1 2026 CALL	85% VERIZON: ISSUES RESOLVED AUTONOMOUSLY, SAME CALL	About 60% T-MOBILE: CHATBOT CONTAINMENT, ONE TRANSCRIPT
34% TELSTRA: ENQUIRIES HANDLED BY AI, FY26 REPORT		

FILED

Verizon call of 27 April 2026, MarketBeat transcript, with RCR Wireless (28 April) and Light Reading (30 April) carrying both figures · Telstra annual report, year to 30 June 2026

REPORTED

T-Mobile call of 28 April 2026, one third-party transcript (MarketBeat)

The most specific statements by an American carrier are Verizon’s of 27 April 2026, and every one of them is missing either a period or a base (Exhibit 4). Its chief executive said Verizon had “over \$200 million of energy savings as a result of deploying AI into the network”, and that “85% of all of our issues right now are autonomously resolved”. He did not say which issues, out of how many, or over what period. Set against Verizon’s own energy total, \$200 million is a large share of what it uses, which is why the missing period matters.

EXHIBIT 4

Verizon's five AI statements of 27 April 2026 give a result or an expectation, and none gives both a period and a base

STATEMENT	RESULT OR EXPECTATION	PERIOD	BASE GIVEN
"over \$200 million of energy savings as a result of deploying AI into the network"	Result: "We already have"	Not stated	None
"85% of all of our issues right now are autonomously resolved"	Result	"right now"	None: which issues, of how many
Vendor-support costs in software development down "by over 70%"	Expectation: "we see our way to"	Not stated	None
Software delivery up "40%+"	Expectation: "opportunities"	Not stated	None
Voice agents: "a 1,280 basis point improvement in customer sat scores"	Result	Year on year	Score not named

All five are in the MarketBeat transcript and in RCR Wireless's report of 28 April 2026; Light Reading's of 30 April carries the first two.

Source: Verizon first-quarter 2026 call, 27 April 2026 (Filed by the series method); the classification of each statement is Blue Prysm's reading of the quoted words (Judgement).

SIZING VERIZON'S \$200 MILLION AGAINST ITS OWN METER

Verizon publishes its energy in gigajoules rather than dollars. Its EY-reviewed schedule gives 41,076,950 GJ for 2025, which is 11.4 TWh, 0.2% below its 2022 figure on the same SASB definition.

Converting \$200 million into energy needs a price, and Verizon does not publish one. The nearest substitute is the value per kilowatt-hour that AT&T's own energy projects saved between 2022 and 2025, which was 8.7 to 11.4 cents ([part 1](#)).

On three assumptions, that the saving is for one year, that all of it is purchased energy and that AT&T's price applies, the \$200 million comes to about 1.75 to 2.3 TWh. That would be 15% to 20% of everything Verizon used in 2025, and more than three times the 4.7% that AT&T's own projects conserved that year.

The call does not say whether the figure covers one year, several years or a saving against forecast growth. A total that was flat from 2022 to 2025 cannot tell those readings apart (Calculated).

The other two figures need the same treatment. T-Mobile's 60% is one third-party transcript's rendering of what its chief financial officer said, with no definition of what counts as contained. Its filed care figure, "more than 50% reduction in calls to care since 2021", runs from three years before its OpenAI platform was

announced in September 2024, so the two cannot be matched. Telstra's 34% is the one AI share in a filing, and it is a share handled rather than a saving. The same report shows labour expenses down 4.6% and underlying operating expenses down A\$454 million, or 3.0%, and attributes neither to AI.

Press figures are weaker still. AT&T's chief financial officer told Fortune in June 2023 that AI had cut employee effort by nearly 17 million minutes a year, worth "hundreds of millions of dollars", and added that the saving was not necessarily related to generative AI. The "up to 90% cost savings" its data chief described to VentureBeat in February 2026 is the cost of running its internal assistant rather than a saving in the operation. The most useful press figure comes from Vodafone. Its group chief technology officer told TM Forum Inform that Vodafone "recognized over €100 million net benefits from AI in the 2025/26 financial year" and had "turned off hundreds of use cases" along the way. The figure is net and comes with a record of stopping things. It does not appear in the 20-F for the year.

The accounts cannot test the 2026 claims until February 2027

A gross saving only shows as a fall in the pool when nothing else is growing, and in 2026 acquisitions and special items are growing faster than any saving could show through. The first half of 2026 shows why the test has to wait (Exhibit 5).

Verizon's SG&A rose 5.9% in the first half of 2026. Inside that line is a \$746 million loss on a business held for sale and \$397 million of new severance, while its Business segment cut personnel costs by \$294 million "related to workforce reduction initiatives". The saving is there but swamped by the rest. T-Mobile's SG&A rose 8.4% against revenue up 9.2%, with UScellular in the base for the whole half. AT&T's other cost of revenues fell 1.4% and its SG&A rose 3.2%.

EXHIBIT 5

In the first half of 2026 the cost lines moved with acquisitions and special items, so no AI saving can be read from them

CARRIER	LINE	CHANGE, H1 2026 ON H1 2025	CONTEXT IN THE 10-Q
Verizon	SG&A	+5.9%	A \$746 million loss on a business held for sale; \$397 million of new severance
Verizon	Business segment personnel costs	-\$294 million	“related to workforce reduction initiatives”
T-Mobile	SG&A	+8.4%	Revenue up 9.2%; UScellular in the base for the whole half
AT&T	Other cost of revenues	-1.4%	“lower personnel and other costs resulting from the decommissioning of our legacy network”
AT&T	SG&A	+3.2%	Lumen’s mass-markets fibre business closed on 2 February 2026

Changes are on the same half of the previous year, as reported.

Source: AT&T, Verizon and T-Mobile Forms 10-Q for the quarter to 30 June 2026 (Filed); percentage changes calculated from the filed lines (Calculated).

Feb 2027

When the Forms 10-K for 2026 arrive, and with them the first filed test of Verizon’s \$5 billion and T-Mobile’s \$1.3 billion. The third-quarter calls in late October 2026 come first, but they are statements rather than accounts.

Until then the only AI operating figure in any of these operators’ annual reports is Telstra’s share of enquiries. None of the eleven has published an AI-specific saving with a period, a base and a filing behind it.

Terms used

- **Run-rate saving.** The annual saving a programme delivers once it is fully in place. A company can reach it before a full year of it shows in the accounts.
- **Core Adjusted EBITDA.** T-Mobile’s earnings measure before interest, tax, depreciation and amortisation and its stated adjustments. A contribution to it can come from cost or from revenue, which is why its \$3 billion is read with care.

- **Containment.** The share of customer contacts an automated channel resolves without a person. No operator here has published its definition.
-
-

Implications

CARRIER STRATEGIST

If you want the AI share of a programme believed, publish it with its period and its base. No operator has yet done so, which means the first to file an AI saving against a named cost line will set the standard the rest are read against.

INVESTOR

Sort every AI figure by where it was said before you model it. Verizon's and T-Mobile's 2026 targets can be tested in the Forms 10-K for 2026. A call figure with no period and no base, such as Verizon's \$200 million of energy, is a statement of direction and nothing more. The three carriers' 10-Ks mention AI 56 times for 2025 and never attach a dollar to it.

VENDOR

Bring a case with a period, a base and a control, because no operator has published what an AI product saved it, and yours may be the first the customer can check. Expect to be asked whether the result is gross or net of the run cost. Vodafone counts net.

Method and limits

HOW THIS WAS BUILT

Each operator's record for 2023 to 2026 was reviewed: annual reports and Forms 10-K, 20-F and 8-K, earnings-call transcripts, company releases and the trade press. Every quantified statement was sorted by kind and tagged by where it was made. A call statement is tagged Filed where two transcripts, or a transcript and a same-week report, give the same figure, and a figure found in only one third-party transcript is tagged Reported. The AI count covers the main document of each 10-K, every section, exhibits excluded. It counts each whole-word "AI", hyphenated forms such as AI-driven included, and each "artificial intelligence" in any case, with "artificial intelligence (AI)" counted once.

WHAT IT DOES NOT SHOW

Third-party transcripts sometimes get names and numbers wrong, which is why single-transcript figures are tagged Reported. A statement absent here is absent from the documents reviewed, which do not include Deutsche Telekom's October 2024 capital markets day, Orange's 2025 registration document or Telefónica's annual report for 2025. Acquisitions (UScellular at T-Mobile, Frontier at Verizon, Lumen's fibre business at AT&T) mean that the cost changes of 2025 and 2026 mix savings with acquired cost, and the two cannot be fully separated.

Data as of: filings, transcripts and reports from 2023 to the second-quarter 2026 results, read 7 Oct 2026 · Method version 1.0.


Sources. AT&T Forms 10-K for 2023 and 2025, 8-K exhibits of 24 January 2024 and 3 December 2024, and call transcripts of 28 January and 22 April 2026; Verizon's 10-K for 2025, 8-K of 20 November 2025, 10-Q for June 2026 and call transcripts of 30 January and 27 April 2026, with RCR Wireless and Light Reading; T-Mobile's results release of 11 February 2026 and call transcripts of 11 February and 28 April 2026; BT's Annual Report 2026 and AP on 18 May 2023; Telstra's FY26 annual report; Vodafone's 20-F for FY2026 and TM Forum Inform; Telefónica's capital markets day deck and exit-plan 6-K; Orange's strategy page; China Mobile's Annual Report 2025; SK Telecom's FY2025 results; Deutsche Telekom's CR report 2025; Fortune on AT&T (29 June 2023); VentureBeat on AT&T (26 February 2026). Read 7 Oct 2026.

Ten AI candidates ranked by the cost of finding out

Which AI candidates can produce a trustworthy number cheapest, and what has been published for each?

Ten AI candidates come up again and again on operators' lists, and the published evidence behind them ranges from a national rollout with a kilowatt-hour total to nothing at all. They are ranked here by one test: the cost of getting a number a finance director would accept. On that test network energy saving comes first. The meter already exists, a matched site makes the control, and a result reads in weeks.

BLUE PRYSM · ANALYSIS



0.63 kWh per site per day

IN BRIEF

- Energy saving ranks first of ten candidates, and it is the only one with a national operator figure. BT's cell sleep on over 19,500 EE sites saves 4.5 GWh a year, which works out at about **0.63 kWh** per site per day against a headline of up to 2 kWh (BT, 24 June 2024).
- 3GPP's own simulations in TR 38.864 put one energy technique's saving anywhere from **0.9% to 84.8%**, so a percentage quoted without its load and its baseline says nothing about what a rollout would save.
- Customer care has results, but each one is for a single journey. Vodafone Portugal lifted first-time resolution for one journey from 15% to 60%, and Telstra's agent assistant takes over a minute off a call. No operator publishes an all-contact containment rate with a definition.
- The largest claims and the thinnest baselines belong to fault correlation, field dispatch and procurement. A carrier should fund energy, agent assist and RAN optimisation first, because each reads within a quarter on data it already holds.

A number costs least to trust where there is a meter and a control

The order here is the order in which a pilot can produce a number that survives a finance review. Ranking the same ten by the size of the prize would give a different list, and the two are often confused.

The logic is borrowed from the [AI opportunity portfolio model](#), with value held equal across the ten. That model divides expected value by the cost of the cheapest test that would settle the question. It rates confidence High only when the effect has been measured “in your own business or in a published result you can check”, and Low when “Somebody believes it”. A published result with a site count, a period and a baseline moves a candidate up that scale. A vendor’s “up to” does not. Each candidate is therefore marked on the five things that set what a test costs (Exhibit 1).

EXHIBIT 1

Scored on five tests, network energy saving earns 9 points of 10 and field dispatch and procurement 3 each

		Data already held	Result in weeks	Control easy to build	Low risk to service	Operator result public	Score of 10
1	Network energy saving	●	●	●	◐	●	9
2	Care: agent assist	●	◐	●	●	◐	8
3	RAN optimisation, SON	●	●	◐	◐	◐	7
4	Software engineering	●	●	◐	●	○	7
5	Churn and offers	●	○	●	●	◐	7
6	Care: deflection	◐	◐	◐	◐	◐	5
7	Fraud and scam	●	○	◐	◐	◐	5
8	Fault correlation, AIOps	◐	○	◐	◐	◐	4
9=	Field-force dispatch	◐	○	◐	◐	○	3
9=	Procurement, contracts	◐	○	○	●	○	3

● yes, 2 ◐ partly, 1 ○ no, 0 Ties: quicker result first, then the published record.

Source: marks are Blue Prysm’s judgement from the published results cited in Exhibit 2: operator and vendor releases, TM Forum case studies, Light Reading and 3GPP TR 38.864, as of 7 October 2026 (Judgement). Note: yes scores two, partly one, no nothing; ties are broken by time to a result and then by the published record.

The marks produce the ranking in Exhibit 2. Energy, agent assist and RAN optimisation all use data the operator already holds, read within a quarter and allow a control. Fault correlation, field dispatch and

procurement carry the largest claims and the thinnest baselines, and rank low for that reason alone.

EXHIBIT 2

Energy saving ranks first of ten and field dispatch and procurement last, because no operator has published a baseline for either

RANK AND CANDIDATE	STRONGEST PUBLISHED RESULT	PUBLISHED BY	TIME TO A RESULT	MEASUREMENT TRAP
1 Network energy saving	4.5 GWh a year across 19,500 sites	BT, operator	Weeks	“Up to” percentages without load or baseline
2 Care: agent assist	20% less follow-up contact (trial); calls over a minute shorter	Telstra, operator	1 to 3 months	Handle time can fall while repeat contacts rise
3 RAN optimisation, SON	Spectral efficiency up to 10% at about 43 sites	T-Mobile with Ericsson, trial	Weeks	The baseline is the vendor’s “rule-based methods”
4 Software engineering	37% of suggestions accepted; seat counts	BT; Verizon	Weeks	Acceptance and seats measure adoption
5 Churn and offers	Churn down 36% (2020); conversion doubled	China Unicom; Safaricom with Huawei	Billing cycles	No holdout design published
6 Care: deflection	Journey automation “approaching 50%”; resolution 15% to 60% on one journey	BT; Vodafone Portugal	6 weeks to 3 months a market	No all-contact rate with a definition
7 Fraud and scam	2.43 million scam calls stopped in four months	BT	Months	Blocked volume is not loss avoided
8 Fault correlation, AIOps	Up to 85% of common faults fixed in one step	Indosat	Quarters	Self-assessed; repair-time baselines move
9= Field-force dispatch	No truck-roll figure; 180,000 callouts prevented, no year	Orange	Quarters	No public baseline
9= Procurement, contracts	£2B of spend through an AI sourcing tool	BT	Quarters	Negotiated is not realised

Source: rank and marks by Blue Prysm from the published results cited in each section below (Judgement); results as published by the operator or vendor named (Reported).

Energy is the cheapest candidate to test

Energy has the shortest path to a number anyone can check, because the meter and the control site both exist before the pilot starts. Nothing has to be built to measure it.

BT holds the one national figure. In June 2024 it switched on cell sleep across “over 19,500 sites across the UK” of its EE network, using “machine learning to predict quiet periods of demand”. The function itself came from “the respective RAN equipment supplier” on each site. The release gives two numbers. One is

“up to 2 KWh per site per day”. The other is “4.5m KWh per year across EE’s estate”. Divided by the sites and the days, the second is about a third of the first.

BT’S CELL SLEEP OF 24 JUNE 2024: THE HEADLINE AND THE AVERAGE

BT reports 4.5 million kWh a year saved across over 19,500 EE sites. That comes to about 231 kWh per site a year, or 0.63 kWh per site per day, a third of the “up to 2 KWh” headline.

The 4.5 GWh is about 0.15% of the 2,989 GWh BT Group used in its 2025 financial year (BT Annual Report 2026, page 71).

The per-site figures and the share of BT’s energy are calculated from BT’s published numbers (Calculated).

Every other published energy figure is a forecast, a trial, an estimate or a vendor’s own statement (Exhibit 3). SoftBank expects about 5 million kWh a year once its sleeping cells grow from about 14,000 to 24,000. Samsung says its AI energy manager on Verizon’s network gave “an energy saving gain of 15% on average, with a maximum of 35% per sector during low traffic periods”. Ericsson says Vodafone UK cut the daily power of its 5G radio units “by up to 33 percent at select sites across London”. Neither vendor says how many sites or for how long. Telefónica’s engineers estimate, without a measured result, that power-saving features save 20% to 30% of RAN energy and that AI might add another 5% to 10% on top.

EXHIBIT 3

Six other energy figures are forecasts, trials, estimates or vendor statements, and the two vendor figures state neither a site count nor a period


WHO	FIGURE	WHAT IT IS	DATE
SoftBank	About 5 million kWh a year expected as sleeping cells rise from about 14,000 to 24,000	Forecast	Jan 2026
Samsung, on Verizon's network	"15% on average, with a maximum of 35% per sector during low traffic periods"	Vendor statement; no site count or period	Feb 2025
Ericsson, with Vodafone UK	5G radio unit daily power cut "by up to 33 percent at select sites across London"	Vendor statement; no site count or period	Mar 2025
Globe, with Nokia	"the equivalent of annual power savings of 3% to 6%"	Proof of concept	Nov 2023
Telefónica	Power-saving features 20% to 30% of RAN energy; AI could add 5% to 10%	Engineers' estimate, no measured result	Sep 2026
Vodafone Türkiye	Up to 10% from software and hardware power saving	Trial; credits "advanced algorithms", not AI	Jul 2026

Source: Light Reading on SoftBank (9 January 2026); Samsung blog (26 February 2025); Ericsson release (11 March 2025); TM Forum Inform on Globe (23 November 2023); Operator Watch on Telefónica (8 September 2026); Vodafone release (8 July 2026) (Reported; vendor figures are the vendor's own).

The standard itself explains why the percentages scatter so widely (Exhibit 4). 3GPP's Release 18 study, [TR 38.864](#), reports ranges from many companies' simulations rather than single values, and the ranges are wide. Nine sources found that lengthening the synchronisation-signal period up to 1,280 ms "could achieve BS energy savings by 0.9%~84.8% in range, meanwhile when traffic occurs and load increases, the UPT significantly decreases". Aligning devices' connected-mode sleep cycles with the cell's saved anywhere from 0.2% to 71.4% across six sources. Three of those found user throughput fell by 0.91% to 15.5%, and one found losses of up to 62.4% in some configurations. Load and the starting configuration move the answer by two orders of magnitude, so a percentage on its own is not a result.

EXHIBIT 4

3GPP's own simulations put one technique's saving anywhere from 0.9% to 84.8%, so a percentage without its load and baseline says nothing about a rollout



Simulated saving in base-station energy, each against its source's own baseline and load.

Source: 3GPP TR 38.864 v18.1.0, Release 18 study on network energy savings, clause 7 conclusions (March 2023), ranges as written there (Filed). Note: each range spans the sources the report counts for it: 8, 9, 5, 6, 3 and 1, in the order shown. The first row is the saving on the sleeping carrier; the anchor carrier used 2.3% to 18.9% more. Throughput costs are in the text.

Two cautions apply to any energy result. The first is about intensity. Light Reading showed that Telefónica's energy per petabyte fell from 268 MWh in 2016 to 54 MWh five years later, while traffic grew more than fourfold. Total consumption fell "by as little as 11%". An intensity figure can collapse while the bill barely moves. The second is about attribution. Vodafone Türkiye's July 2026 trial, up to 10% from software and hardware power saving, credits "advanced algorithms" rather than AI, and there is no reason to relabel it.

Energy ranks first because a test is cheap and the unit is physical. It is also the candidate where headline and measurement lie furthest apart. BT's "up to" figure is three times its own average, and the standard's simulations span two orders of magnitude. The pilot that settles it counts kilowatt-hours at matched sites and measures user throughput alongside, so that a saving bought with dropped sessions shows up as what it is.

Care has results by journey but no all-contact rate

Customer care is where most operators’ published AI results land, and every one of them is for one journey, one market or one tool (Exhibit 5).

BT says its assistant Aimee handles about 60,000 conversations a week, double what it handled two years earlier, with “automation success rates on several types of customer journey now approaching 50%”. Vodafone’s SuperTOBi lifted first-time resolution for appointment booking in Portugal “from 15% to 60%”, though its early answers were right only 75% to 78% of the time before rework took them to about 90%. Its chief technology officer said a less cautious operator could launch in “six weeks to three months”. Telstra’s agent assistant, in the hands of more than 8,000 staff, is “cutting average customer call time by over a minute”. These results are real but narrow.

EXHIBIT 5				
Operators’ care results are for one journey, one market or one tool, and none is an all-contact containment rate with a definition				
OPERATOR	TOOL	RESULT	SCOPE	DATE
BT	Aimee, virtual assistant	About 60,000 conversations a week, double two years earlier; automation success “approaching 50%”	Several journey types	Dec 2024
Vodafone	SuperTOBi, Portugal	First-time resolution for appointment booking “from 15% to 60%”; answers right 75% to 78%, then about 90% after rework	One journey, one market	Jul 2024
Telstra	Call-summary tool	Follow-up contact down 20%	Trial, 2023	Feb 2024
Telstra	Agent assistant, 8,000+ staff	“cutting average customer call time by over a minute”	Agent tool	Apr 2026
Telstra	AI Assistant	Handled 34% of enquiries since March 2026	All enquiries; a share handled	FY26 report
DNA, Finland	Process redesigned around AI	“reduced customer care work by 34%”	One process	Sep 2026
Source: BT newsroom (12 December 2024); Vodafone release (4 July 2024) and Light Reading (5 July 2024); Telstra releases (7 February 2024, 20 April 2026); TM Forum Inform (22 September 2026) (Reported); Telstra Financial Results and Annual Report FY26 (Filed).				

Two documents mark the limits. Verizon’s own survey of 5,000 consumers in seven countries found “88% of consumers are satisfied with interactions handled mostly or fully by human agents”, against 60% for AI-driven ones. McKinsey described an agent copilot that “realized just a 5 percent improvement in productivity” until 90% of its budget was moved to training and change management, after which the gain passed 30%. Agent assist therefore ranks above deflection: a person stays in the loop, and the test splits cleanly by agent group.

A deflection pilot that reports containment without repeat contacts and complaints beside it is reporting the one measure that rises whether or not the customer's problem was solved. No operator in the published record states an all-contact containment rate with its definition, and until one does, every containment figure is a number without a denominator.

The network-operations claims are large and mostly self-assessed

The fault-correlation and dispatch figures are the largest in the published record and the hardest to test. The data needs cleaning before a pilot can start, results take quarters rather than weeks, and the baseline moves as the process changes under it.

- **Fault correlation.** Indosat reports resolving “up to 85% of common wireless access faults in a single step” and cutting raw alarms “by up to 95%” across more than 55,000 base stations. It also says that “wrangling and cleansing the data is 70% to 80% of the time and effort”. It publishes no repair time and no ticket count. China Mobile's pilots with Huawei report average repair time down 30%, on its own assessment.
- **Automated change.** Verizon made “more than 70 million network configuration changes autonomously” in 2025. The figure measures activity rather than savings.
- **RAN optimisation.** T-Mobile and Ericsson ran AI link adaptation on live traffic at about 43 sites and reported spectral efficiency up to 10% and throughput up to 15% against “rule-based methods”. Orange changes 20 million radio parameters a month automatically. RAN optimisation ranks third because a cluster trial reads in weeks on counters the operator already holds.
- **Field dispatch.** No operator in the published record reports a fall in truck rolls with a sample behind it. Orange says AI diagnostics “prevented 180,000 maintenance callouts in France”, with no year and no base. Indosat's field copilot raises “field productivity by 20% to 30%”, by its own account.

REPORTED

TM Forum on Indosat (3 Sep 2026) and China Mobile with Huawei (3 Sep 2025); Light Reading on Verizon (24 Jun 2026), T-Mobile and Ericsson (22 May 2026) and Orange (15 Feb 2024); Orange AI page

The rest of the list

- **Fraud and scam.** BT stopped more than 2,430,000 scam calls in four months for its 2.5 million Digital Voice customers, which is about one call each. Bell Canada's machine-learning risk score gave a “10% increase in fraud detection”. A blocked call is a volume, and a volume is not a loss avoided.
- **Churn and offers.** China Unicom reported churn down 36% in 2020. Safaricom, with Huawei, doubled conversion on targeted packages for about 3 million subscribers. Neither published a control, so neither result can be separated from whatever else was happening to those customers at the time.

- **Software engineering.** BT's first four months with CodeWhisperer produced over 100,000 lines of code with 37% of suggestions accepted. Verizon took Claude Code from fewer than 500 users to 33,000 in about six weeks. Both figures measure adoption, and neither operator has published what the code did for it.

- **Procurement.** BT put “around £2bn of spend” through an AI sourcing platform and cut simple purchases “from days and weeks to seven minutes”. Process time is published; realised savings are not.

REPORTED

BT newsroom on scam calls (3 Oct 2024) and CodeWhisperer (20 Feb 2024); TM Forum on Bell Canada (20 Oct 2021), China Unicom (23 Jun 2020), Safaricom (22 Oct 2025) and BT Sourced (6 Apr 2023); Light Reading on Verizon (24 Jun 2026)

One result belongs beside the ranking rather than in it. Vodafone's net benefit of over €100 million in FY26, which [part 2](#) describes, came with “hundreds of use cases” turned off along the way. A portfolio ranked by the cost of finding out is built to stop the same use cases earlier, before they have cost much. [Part 5](#) sets out the rules for doing that.

Terms used

- **Cell sleep.** Switching off a carrier, a cell or parts of a radio when traffic is low. 3GPP's Release 18 and 19 features let more of a cell sleep; the decision about when is the vendor's or the operator's.
- **Containment.** The share of contacts an automated channel resolves without a person.
- **First-time resolution.** The share of problems solved in the first contact.
- **SON.** Self-organising networks: software that tunes radio parameters automatically.
- **Holdout.** Customers kept out of a treatment so that its effect can be measured against them.

Implications

CARRIER STRATEGIST

Fund the top three first. Energy, agent assist and RAN optimisation each use data you already hold, read within a quarter and allow a control. A measured result from any one of them is worth more to the programme than a forecast from all the rest.

INVESTOR

Discount an AI saving by its row in Exhibit 2. A figure from the top of the table can be checked against a meter or a control. A figure from the bottom has no published baseline at any operator in the record, so a large number there is a forecast, however confidently it is stated.

VENDOR

Publish site counts, periods, baselines and controls with your percentages. The operator's confidence rating, and with it your place in its portfolio, rests on a result it can check for itself, and an "up to" figure scores as asserted and nothing more.

Method and limits

HOW THIS WAS BUILT

Each candidate is marked on five things that set the cost of a trustworthy test. The first three are whether the operator already holds the data, whether a result can be read in weeks and whether a control is easy to build. The other two are whether the test is low-risk to live service and whether an operator has published a result. Yes scores two, partly one and no nothing. The rank is the total, with ties broken by time to a result and then by the published record. The marks are Blue Prysm's judgement. Figures are as published by the operator or vendor named, and a figure that comes from a vendor alone is marked as such.

WHAT IT DOES NOT SHOW

The ranking is by the cost of finding out, which is not the same as value. A candidate low on the list may be worth more than one at the top once it has been tested, and the [portfolio model](#) puts value back in. None of the figures here was independently measured, and many of the TM Forum case studies were written with a vendor at the table. Absence here means absence from the published record reviewed; an operator may hold results it has chosen not to publish.

Data as of: operator, vendor and trade-press publications from 2020 to September 2026 and 3GPP TR 38.864 v18.1.0, read 7 Oct 2026 · Method version 1.0.

Sources. Energy: BT on cell sleep (24 June 2024); Light Reading on SoftBank and Telefónica; Samsung and Ericsson; TM Forum on Globe; Operator Watch on Telefónica; Vodafone Türkiye; 3GPP TR 38.864. Care: BT on Aimee; Vodafone on SuperTOBi and Light Reading on its rebuild; Telstra's February 2024 release and April 2026 update; TM Forum Inform on DNA and Vodafone; Verizon's customer-experience survey; McKinsey. Operations: TM Forum on Indosat and China Mobile; Light Reading on Verizon, T-Mobile and Orange; Orange AI. The rest: BT on scam calls and CodeWhisperer; TM Forum on Bell Canada, China Unicom, Safaricom and BT Sourced. Read 7 Oct 2026.

What the standards add

What do 3GPP, O-RAN and the autonomy levels give an operator to build an AI programme on?

What the standards give an operator to build an AI programme on has two halves. 3GPP has specified where a network's analytics data comes from, how a model is trained, monitored and rolled back, and what each autonomy level means. It has barely specified anything that lets one vendor's model work inside another vendor's network. Every result that would show any of this paying is held by an operator or a vendor rather than a standards body.

BLUE PRYSM · ANALYSIS

3.42 average · L4 in 21 scenarios

IN BRIEF

- NWDAF has been in the 5G core since Release 15, but a standard way to monitor a model's accuracy arrived only with TS 23.288 v18.1.0 in March 2023, almost four years after the first full version. Release 17 limited model sharing to **single-vendor** environments.
- The radio study TR 38.843 (December 2023) found no consensus for two-sided CSI compression. Release 19 standardised one-sided models, implementable from December 2025, which is why near-term AI in the RAN comes from the incumbent RAN vendor.
- One third-party rApp runs in a live network, AT&T's since **30 July 2025**, with no published result. Omdia found live near-real-time RIC deployments "limited", and Ericsson says the RAN itself already does what xApps would.
- Autonomy levels are self-assessed. China Mobile targeted L4 for almost all services by 2025 and now reports L4 in **21** scenarios; TM Forum's survey found 21% of companies at L3 or above. A buyer should ask for the scope behind every level before taking it seriously.

Every 3GPP release since 2019 has added something an AI programme can use

Three groups inside 3GPP write the parts of the standard an AI programme uses. The core-network groups define where analytics data comes from and how a model may be shared. The management group defines how a model is trained, tested, deployed, monitored and rolled back. The radio groups define where a model may live in the air interface, and the signals that let a cell save energy (Exhibit 1).

EXHIBIT 1

Release 18 brought accuracy monitoring and the full model life cycle; Release 19 brought one-sided radio models, degradation thresholds and fallback actions

RELEASE	CORE NETWORK	MANAGEMENT	RADIO
Rel-15	NWDAF specified in the first 5G release		
Rel-16 (TS 23.288 v16.0.0, June 2019)	Data collected from network functions, applications and OAM; outputs as statistics and predictions		
Rel-17 (v17.0.0, March 2021)	Training and inference split; model sharing “limited to single vendor environments”	Analytics for coverage, faults and energy saving (TS 28.104, June 2022); ML training; autonomy levels	
Rel-18 (v18.0.0, December 2022)	Accuracy monitoring of models and analytics, from v18.1.0 (March 2023); federated learning	Full AI/ML life cycle (TS 28.105); energy-saving predictions	AI/ML study TR 38.843; AI/ML for energy saving, load balancing and mobility; energy-saving signalling
Rel-19 (implementable from December 2025)	Vertical federated learning; QoS and policy from predicted experience	Learning on the live network with degradation thresholds and fallback actions ; software upgrade validation	One-sided models for beam management, positioning and CSI prediction; on-demand SSB and SIB1
Rel-20 (in progress)			Two-sided models for CSI compression

Dates are TS 23.288 version dates from the 3GPP archive. Release 19’s date is 3GPP’s own statement that its specifications were “fully implementable” at the end of December 2025.

Source: 3GPP TR 21.916, 21.917, 21.918 and 21.919; TS 23.288, 28.104 and 28.105 archives; 3GPP NWDAF and Release 19 papers (Filed).

The core got its analytics function in 2019 and a way to check it in 2023


The network data analytics function, NWDAF, was “specified from the 5G’s initial Release (Rel-15)”, in 3GPP’s own words, and its specification, TS 23.288, reached version 16.0.0 in June 2019. Release 16 said what it collects and what it puts out. Release 17 split model training from inference, added functions to

coordinate data collection and store analytics, and allowed “Trained data model sharing between multiple NWDAF instances, limited to single vendor environments”. Release 18 added federated learning in version 18.0.0 of December 2022 and then, in version 18.1.0 of March 2023, the means to compute “the accuracy of the ML models and analytics”. Release 19 added vertical federated learning and policy recommendations driven by predicted quality of experience.

Two of those dates matter more than the others (Exhibit 2). The first is March 2023, because a standard way to check whether a model in the core is still accurate arrived almost four years after the specification’s first full version. Until then an operator could run a model but had no standard way of knowing whether it had drifted. The second is March 2021, when sharing a trained model between vendors was written out of Release 17. Together they mean that a multi-vendor analytics estate built on NWDAF is recent territory for the standard, and no operator in the published record has reported a result from a commercial NWDAF deployment.

EXHIBIT 2

NWDAF reached its first full specification in June 2019; a standard way to check a model’s accuracy followed in March 2023 and live-learning controls in December 2025



Source: 3GPP specification archive for TS 23.288, TS 28.104, TS 28.105 and TS 28.100; 3GPP releases page for the Release 18 functional freeze and the Release 19 page for “fully implementable” (Filed); AT&T release of 30 July 2025 (Reported). Note: dates are the archive dates of each version.

The radio study recommended the uses one vendor can build alone

The air-interface study, [TR 38.843](#), was approved in December 2023, and its conclusions matter as much for what they decline as for what they recommend. On compressing channel-state information with a two-sided model: “From RAN1 perspective, there is no consensus on the recommendation of CSI compression

for normative work”. The reasons given are the “Trade-off between performance and complexity/overhead” and “Issues related to inter-vendor training collaboration”. On predicting channel state there was no consensus either, for “Lack of results on the performance gain over non-AI/ML based approach and associated complexity”. What the study did recommend for normative work was beam prediction and positioning, both of which one vendor can build at one end of the link.


Release 19 then gave “normative support for a general framework of one-sided (at the gNB or UE sides) AI/ML models as well as ... beam management, positioning enhancements and CSI prediction”. Offline training is assumed for positioning. The two-sided model went to a Release 20 work item instead (Exhibit 3). [6G part 8](#) covers that alongside the rest of the AI-native argument.

EXHIBIT 3

The standard has adopted the AI a single vendor can implement at one end of the link and deferred the AI that needs two vendors to agree


AI in the radio: one-sided

A model lives wholly in the phone or wholly in the network.
Rel-19 normative: beam prediction, positioning.
Everyone agrees.




AI across the link: two-sided

Half the model in the phone, half in the base station,
trained together. Rel-20 work item; three interoperability
routes still open. Qualcomm for; Ericsson cool.



AI as the radio's computer: AI-RAN

The base station's software runs on general-purpose GPUs.
Not a 3GPP matter. Nokia, NVIDIA and SoftBank for;
Ericsson, Samsung and most operators against, on cost.



Decided (RP-251881): the 6G System can also operate without AI/ML

Source: 3GPP RP-240774 (Release 19), RP-251870 (Release 20) and RP-251881 (6G study); vendor and operator positions as cited in 6G part 8 of this site's analysis (Filed and Reported as marked there). Note: the one-sided row is what Release 19 specified; the two-sided row is the Release 20 work item; AI-RAN compute is outside 3GPP.

For operations the relevant radio items are smaller and more useful. Release 18 specified AI/ML in the radio access network for energy saving, load balancing and mobility, with an “Energy Cost” index that base stations exchange with one another. Release 19 extended the work to slicing and to coverage and capacity optimisation. The energy work adds signalling rather than intelligence, and the distinction matters. Release 18 specified cell discontinuous transmission and reception, secondary cells without synchronisation blocks and the barring of older devices. Release 19 added synchronisation and system information on demand, woken by a device's uplink signal. Each of those is a hook that lets a scheduler put more of a cell to sleep.

The decision about when to sleep, which is where the AI lives, stays with the vendor or the operator, which is why [part 3](#) found its energy results in company announcements rather than in specifications.

The radio standard has adopted the AI uses one vendor can implement alone, at one end of the link, and sent the uses that need two vendors to agree to a later release. The near-term AI value in the RAN therefore arrives as software bought from the operator's existing RAN vendor.

The management plane now names the controls a kill gate needs

The management specifications are where the standard meets the kill gate. Management data analytics, TS 28.104, was first published for Release 17 in June 2022 and covers coverage, fault and energy-saving analytics. The AI/ML management specification, TS 28.105, covered training in Release 17 and grew to the whole life cycle in Release 18: training, testing, emulation, deployment and inference. Release 19 adds reinforcement learning, and with it a set of named controls (Exhibit 4).

The autonomy levels belong to the same family of documents. [TS 28.100](#) defines six levels, from L0, manual, to L5, full, across five kinds of task. L4 is the level at which "All the execution, awareness, analysis and decision tasks are accomplished automatically by telecom system itself. And intent handling tasks can be partly accomplished automatically". Its first note gives the person the final say when a person and the system disagree.

EXHIBIT 4

Release 19 management specifications name four controls for a model that learns on a live network, and TS 28.100 gives a human the final say

- **Degradation thresholds.** TS 28.105 Release 19 adds “tolerable degradation thresholds” for a model that learns on the live network: how far it may slip before something happens.
- **Fallback actions.** The same release adds “fallback actions”, which is what happens when a threshold is crossed.
- **Environment type.** A reinforcement-learning model has to declare its “RL environment type (simulation or real network)”, so that one is not mistaken for the other.
- **Software upgrade validation.** TS 28.104 Release 19 adds “Software upgrade validation” to management data analytics.
- **Human authority.** TS 28.100 note 1: “Human reviewed decision have the highest authority in each level if there is any confliction between human reviewed decision and telecom system generated decision”.

Source: 3GPP TR 21.919 v19.0.0 for the TS 28.104 and TS 28.105 Release 19 content; TS 28.100 V19.0.0 (September 2025), note 1 (Filed).

The standard now writes down what a kill gate needs for a model that learns on a live network. It names a degradation that counts as tolerable, an action to fall back to, and a human decision that overrides the system at every level. None of that is new thinking. [Part 5](#) shows why each of those controls was needed long before AI came into it.

The one live third-party rApp has no published result

O-RAN’s specification set reached 147 unique titles by February 2026, 71 of them new or updated since November 2025. It includes an A1 service for managing ML models and conformance tests for rApps. The adoption record is short by comparison (Exhibit 5). AT&T announced on 30 July 2025 that it had become “the first CSP globally to deploy a third-party rApp to optimize its live production network”, on Ericsson’s automation platform. It did not name the developer and it has not published a result. Every other entry in Exhibit 5 is a validation, a test environment or a vendor’s own figure.

EXHIBIT 5

The open-RAN controller record is one third-party rApp live, three in validation and no market in xApps

WHAT	WHO	STATUS	RESULT PUBLISHED	DATE
Third-party rApp in a live production network	AT&T, on Ericsson's automation platform	Live; developer not named	None	30 Jul 2025
Three Ericsson rApps	Swisscom	"in validation"	None	23 Feb 2026
rApps on Ericsson's platform	Ericsson	Nearly 90, 25 of them Ericsson's own	n/a	23 Feb 2026
rApp generated by AI for AT&T	AT&T with Ericsson	Test environment, not the live network	None	Feb 2026
AI energy manager on a RIC	Samsung at Verizon; TelecomLead reports Qualcomm's RIC	Deployed	Vendor's own: 15% on average, 35% maximum per sector	Feb 2025
Near-real-time RIC and xApps	Ericsson; Samsung; Omdia	Ericsson: "The functionality of xApps is provided by the RAN itself"; Omdia: "To date, live deployments are limited"	None	Dec 2024; Feb 2025
O-RAN specification set	O-RAN Alliance	147 unique titles, 71 new or updated since November 2025; A1 service for ML models; rApp conformance tests	n/a	27 Feb 2026

Source: AT&T release (30 July 2025); Ericsson release on Swisscom (23 February 2026); Telecom Reseller on the AI-generated rApp (February 2026); Samsung blog (26 February 2025) and TelecomLead; Light Reading on xApps (5 December 2024) and rApp marketplaces (4 February 2025); O-RAN Alliance publication note (27 February 2026) (Reported).


The near-real-time controller and its xApps have done worse. Ericsson told Light Reading in December 2024 that the RAN itself provides what xApps would, and Samsung saw no near-term path for a standalone near-real-time RIC. Omdia found live deployments limited in February 2025. The clearest RIC-hosted use with an operator's name on a number is Samsung's energy manager at Verizon, which TelecomLead reports runs on Qualcomm's RIC. [Part 3](#) marks its figures as the vendor's own. In our view an operator planning AI in the RAN is buying from its RAN vendor, whatever the architecture diagram shows.

Autonomy levels are self-assessed

Autonomy levels are the industry's scoreboard for operational AI, and every score on the board was awarded by the operator to itself (Exhibit 6).

EXHIBIT 6

China Mobile targeted L4 for almost all services by 2025 and reports L4 in 21 scenarios in 2026; Telefónica reports a group average of 3.42



Filled: reported. Hollow: target. Band: the range the operator states. All self-assessed.

Source: 3GPP TS 28.100 V19.0.0 for the levels (Filed). China Mobile: TM Forum Inform (10 August 2023) and its interim chairman's statement (13 August 2026); Telefónica: Operator Watch (8 September 2026); Verizon: Light Reading (24 June 2026) (Reported). Note: filled is reported, hollow a target, the band a stated range; every level is self-assessed.

China Mobile set the boldest target. Its Lingli Deng told TM Forum Inform in August 2023: "We are targeting Level 4 by 2025 for almost all the services that we are currently providing", at a time when its production average was 2.85. Its 2024 results said the network "has already reached an advanced level (L4) of smart autonomy in some scenarios". The interim chairman's statement of 13 August 2026 says "Autonomous networks achieved L4-level capabilities in 21 high-value scenarios", with 38 planned over three years. The unit changed along the way: the target was written in services, the result comes back in scenarios, and none of the documents reviewed gives an average. Telefónica reports a group average of 3.42 at the end of 2025, up from about 1.1 five years earlier, with targets of 3.75 by 2028 and Level 4 by 2030. Verizon says it is working beyond Levels 2 and 3 and aims at L4 "in the near future".

The surveys put the claims in proportion. TM Forum's survey of March 2026, which asked 125 people at 80 companies, found 21% at L3 or above overall, against 19% a year before. Bain's survey of 22 operators with TM Forum found 20% at L4 or L5 "in select domains". The AI-RAN Alliance had 132 members by February 2026 and has published reference architectures rather than measured results. The GSMA's benchmarks of language models on telecom tasks found that performance on turning an intent into a configuration "remains low, with top models scoring under 30". That task is the one closest to an AI changing a network by itself.

China Mobile's claim moved from L4 for almost all services by 2025 to L4 in 21 high-value scenarios in 2026, and nothing in the standard stopped it. TS 28.100 defines the level and says nothing about who may award it. An autonomy level is a self-assessment of a named scope and should be read with the scope beside it.

Terms used

- **NWDAF.** Network data analytics function: the 5G core's function for collecting data and producing statistics and predictions for other network functions.
- **MDA.** Management data analytics: the same idea in the management system, for coverage, faults and energy.
- **One-sided and two-sided model.** A one-sided model lives wholly in the phone or wholly in the network. A two-sided model is split across the link, so two vendors have to train it together.
- **RIC, rApp, xApp.** The RAN intelligent controller, in non-real-time and near-real-time forms; rApps run on the first, xApps on the second.
- **L4.** In TS 28.100, the level at which the system performs all execution, awareness, analysis and decision tasks itself.

Implications

CARRIER STRATEGIST

Build on what is specified and implementable now. That means NWDAF with Release 18 accuracy monitoring, one-sided radio models, the Release 18 and 19 energy signalling, and TS 28.105 thresholds and fallback for anything that learns on the live network. Treat cross-vendor model sharing and third-party xApps as roadmap, because neither has a live deployment with a result behind it.

INVESTOR

Expect no revenue line and little opening for new vendors from these standards. The AI in the RAN is bought from the incumbent, and the whole open-market record is one third-party rApp in a live network with no published result.

VENDOR

Ship TS 28.105 degradation thresholds and fallback actions, and accuracy monitoring in the core, because a customer can now ask for them by name. State the scope behind any autonomy level you claim; a level without one will be read as a self-assessment, because that is what it is.

Method and limits

HOW THIS WAS BUILT

The 3GPP content is taken from the specifications themselves and from 3GPP's own release descriptions, TR 21.916 to TR 21.919, with version dates from the specification archive. Where the wording matters it is quoted. O-RAN's position comes from the Alliance's publication note and from operator and vendor releases, and only what an operator has announced as live counts as adoption. The autonomy levels are defined from TS 28.100, and operator claims and targets are quoted as published, without adjustment.

WHAT IT DOES NOT SHOW

A specification says what may be built rather than what has been built. NWDAF, MDA and RIC deployments are counted only from what operators announce, and an operator may run more than it announces. Autonomy levels are self-assessed, often with a vendor as co-author, and cannot be compared between operators. The [6G series](#) covers the AI-native design of 6G, so it is not repeated here.

Data as of: 3GPP specifications and release descriptions to TR 21.919 v19.0.0, O-RAN and operator publications to September 2026, read 7 Oct 2026 · Method version 1.0.


Sources. 3GPP release descriptions TR 21.916 to TR 21.919; the specification archives for TS 23.288, TS 28.104, TS 28.105, TS 28.100 V19.0.0, TR 38.843 and TR 38.864; 3GPP's NWDAF page, AI/ML for NR page and Release 19 page; the O-RAN Alliance's publication note of 27 February 2026; AT&T on its third-party rApp (30 July 2025); Ericsson on Swisscom (23 February 2026); Telecom Reseller on the AI-generated rApp (19 February 2026); Light Reading on xApps, rApp marketplaces and Verizon; TelecomLead on Verizon's energy RIC; TM Forum Inform on China Mobile's target (10 August 2023); China Mobile's 2024 annual results and 2026 interim chairman's statement; Developing Telecoms on DTW 2026; Operator Watch on Telefónica; Fierce on the TM Forum survey; Bain with TM Forum; the AI-RAN Alliance (26 February 2026); the GSMA on its telecom model benchmarks. Read 7 Oct 2026.

Where the kill gate goes

Which rules and which failures decide where an AI pilot must stop, and what should it measure while it runs?

The public record holds no regulator's fine against an operator for its own AI. It holds instead two failures of automated change. At Rogers a risk algorithm rated a national core change Low. At AT&T a network element was loaded without the review the operator's own rules required. Those are the precedents. The rules that bind today are about voices and disclosure, and the European duties for AI in critical digital infrastructure now start on 2 December 2027.

BLUE PRYSM · ANALYSIS



3 minutes from change to outage

IN BRIEF

- In July 2022 Rogers' risk algorithm rated phase six of a core change Low, and over **12 million** customers lost service for about 26 hours. The FCC's report on AT&T's outage of 22 February 2024, which blocked more than 92 million calls, says that loading a change which fails its review criteria "should not be possible".
- Three rules bind now. FCC 24-17 has required prior express consent for AI voices since February 2024, and Article 50(1) of the AI Act has required chatbots to say they are AI since 2 August 2026. Regulation 2026/1744 moves the high-risk duties for critical infrastructure to **2 December 2027**.
- AI built solely for efficiency, which covers energy, RAN optimisation and AIOps, is outside the EU's safety-function test under recital 7. Automated change whose failure could cut emergency calling stays inside the failure test.

Three of the four big outages were changes a control should have stopped

None of the network failures in the public record was caused by AI. Three of them were changes that a control should have stopped, and at Rogers the control that failed was itself a piece of automation (Exhibit 1).

EXHIBIT 1			
Four outages since 2022 cut off tens of millions of customers, and in each one where a cause is published a change got past a control			
OUTAGE	SCALE	WHAT THE CONTROL MISSED	RECORD
Rogers, 8 Jul 2022	Over 12 million customers, about 26 hours	A risk algorithm rated phase 6 of a 7-phase core change Low, down from High; scrutiny and lab tests were skipped	Xona Partners for the CRTC
AT&T, 22 Feb 2024	Over 125 million devices; over 92 million calls and over 25,000 calls to 911 blocked; at least 12 hours	An element was loaded without the peer review AT&T's procedures required; an automated protection mode then shut connections down	FCC report; referred to the Enforcement Bureau
Optus, 18 Sep 2025	About 14 hours; 455 Triple Zero calls failed; two deaths	The contractor used the wrong 2022 procedure and classed the job no-impact; Optus classed it urgent, skipping review; the post-change check was "not sufficiently granular"	Independent review; ACMA court action
Verizon, 14 Jan 2026	Up to 10 hours; 1.5 million outage reports	A "software issue"; no root cause published	Press; no FCC report yet
Source: Xona Partners executive summary, published by the CRTC (4 July 2024); FCC DOC-404150A1 (22 July 2024) (Filed). The Register (19 December 2025) and The Nightly (30 July 2026) on Optus; Light Reading on Verizon (Reported).			

ROGERS, 8 JULY 2022

Xona Partners’ assessment for the Canadian regulator describes the sixth phase of a seven-phase core upgrade. At that point “the risk assessment algorithm downgraded the risk level for the sixth phase of the configuration change to ‘Low’ risk”. It had been rated High.

That downgrade, the report says, “contravenes industry norms”. At Low, the phase skipped the scrutiny and the laboratory testing a High rating would have brought. Over 900,000 routes were released where about 10,000 would be normal, and over 12 million customers lost service for about 26 hours.

The report recommends fewer changes per maintenance window, lab tests that mirror production and “tools and processes for automatic rollback of configuration parameters”.

AT&T, 22 FEBRUARY 2024

A change loaded at 2:42 AM Central time took the network down at 2:45. Rolling it back took “close to two hours”, and restoring service took at least 12 hours. More than 125 million devices lost service, and more than 92 million calls and more than 25,000 calls to 911 were blocked.

The Bureau found “a lack of peer review, a failure to adequately test after installation, inadequate laboratory testing, insufficient safeguards and controls to ensure approval of changes affecting the core network” among the causes.


It then wrote the rule in one sentence: “It should not be possible to load changes that fail to meet those criteria”. The matter was referred to the Enforcement Bureau.

Optus adds the check that comes after the change. Its 2:40 AM check missed rising call failures because “This data was not sufficiently granular to enable detection of the emerging problem”. Its board has since accepted all 21 recommendations of the independent review and expanded its automated testing of Triple Zero calls.

These cases are not an argument against automation. Comcast says “99.7% of the software changes that we are making are fully automated”, and Verizon made more than 70 million configuration changes autonomously in 2025. Automation at that scale is normal. The failure in each case is a control that a change, or a tool, could get past (Exhibit 2).

EXHIBIT 2

The Rogers, AT&T and Optus outages each got past a different step of the change process, and the controls now called for cover the steps an AI would occupy



Source: Xona Partners assessment for the CRTC (4 July 2024); FCC PSHSB report on the 22 February 2024 AT&T outage (22 July 2024); the Optus independent review as reported by The Register (19 December 2025); 3GPP TS 28.105 Release 19 per TR 21.919 (Filed and Reported as marked). Note: a cross marks a step the record says failed; a dot marks a control the record now calls for. Compiled by Blue Prism.

An AI that scores, approves or executes network changes would occupy the steps where these failures happened. The first kill gate therefore goes in change control, in one sentence: no automated tool, AI or otherwise, may lower a change's risk class or load a change that has not passed review.

The rules in force today are about voices and disclosure

The rules that bind an operator's AI today are about voices and disclosure. The duties for AI that runs critical infrastructure are written and dated, and the date has already moved once (Exhibit 3).

United States: voices. In February 2024 the FCC ruled in [FCC 24-17](#) that AI-generated voices “are ‘artificial’ voice messages because a person is not speaking them”. It follows that “callers that use such technologies must obtain the prior express consent of the called party”. That reaches an outbound AI agent making renewal, collections or retention calls. Inbound AI care is not a TCPA call, so it does not. In August 2024 the FCC proposed that AI-generated calls disclose it “at the beginning of each call”, and the trackers reported no final rule as of 5 October 2026.

The FCC’s enforcement so far has been against traffic rather than operators. It levied a \$6 million forfeiture for the cloned-voice Biden robocalls of January 2024. CyberScoop reported a \$1 million consent decree with Lingo Telecom, the carrier that passed those calls with A-level caller-ID attestation.

United States: preemption. Federal preemption of state AI law has been tried and has not happened. The Senate struck a moratorium out of the 2025 budget bill by 99 votes to 1. Executive Order 14365 of 11 December 2025 directed the FCC to “initiate a proceeding to determine whether to adopt a Federal reporting and disclosure standard for AI models”. A tracker reported that proceeding not yet opened as of 5 October 2026. The FTC proposed on 1 July 2026 that a law like Colorado’s is “impliedly preempted to the extent it conflicts with a federal regulatory scheme”. The FTC’s statement is a proposal and not yet a decision.


Colorado. Colorado pushed its AI Act back to 30 June 2026 in SB25B-004, then saw enforcement held in April 2026 after a challenge the Justice Department joined. The law firms Seyfarth and McDermott report that SB 26-189, which replaces it, was signed on 14 May 2026 and takes effect on 1 January 2027. It brings notice duties for automated decision-making in consequential decisions. The legislature’s own bill page still listed only the bill’s introduction, on 1 May 2026. For an operator the law reaches credit and eligibility decisions in sales rather than the network.

European Union. Annex III of the AI Act lists “AI systems intended to be used as safety components in the management and operation of critical digital infrastructure”, and that takes in “Providers of public electronic communications networks”. Regulation (EU) 2026/1744 of 8 July 2026, in force from 27 July, moved the high-risk duties for Annex III systems to 2 December 2027 and narrowed what counts as a safety function. Its recital 7 says the term “does not include AI systems which are intended to solely fulfil functions related to user assistance, performance optimisation, service efficiency, automation, convenience”. The amended definition keeps the second test. A safety component is still one “the failure or malfunctioning of which endangers the health and safety of persons or property”.

The chatbot duty did not move. Under Article 50(1) a system that talks to people has to tell them it is an AI “unless this is obvious”, and that duty has applied since 2 August 2026. Most breaches can be fined up to €15 million or 3% of worldwide turnover. The NIST AI Risk Management Framework is voluntary and under revision.

EXHIBIT 3

The rules binding an operator's AI arrived between February 2024 and August 2026; the EU's infrastructure duties follow in December 2027



Source: FCC 24-17 (8 February 2024), 24-84 (Federal Register, 10 September 2024) and 24-104 (30 September 2024); Executive Order 14365 (11 December 2025); FTC release (1 July 2026); Regulations (EU) 2024/1689 and 2026/1744 (Filed). Wisconsin Examiner on the Senate vote; Seyfarth and McDermott on Colorado; Governome and The Political Group on the FCC items (Reported). Note: the Colorado dates are as reported by law firms.

Each of those rules reaches a different candidate from [part 3](#) (Exhibit 4). The narrowed EU definition takes AI built solely for efficiency out of the safety-function test altogether. It leaves the failure test, which catches automated change that could cut emergency calling if it went wrong: the change the AT&T, Rogers and Optus records describe. An operator therefore has until 2 December 2027 for that class of system, and no time at all for chatbot disclosure.

EXHIBIT 4

Consent and disclosure duties bind now; the EU's high-risk duties for automated change that could endanger safety start on 2 December 2027

CANDIDATE	RULE	FROM	WHAT IT REQUIRES
Outbound AI voice: retention, collections, renewals	TCPA, as read in FCC 24-17	8 Feb 2024	Prior express consent of the person called
Chatbots and voice agents in the EU	AI Act, Article 50(1)	2 Aug 2026	Tell people they are dealing with an AI, unless obvious
Automated network change whose failure could endanger safety, such as emergency calling, in the EU	AI Act, Annex III point 2, as amended	2 Dec 2027	High-risk duties
Energy saving, RAN optimisation, AIOps built for efficiency, in the EU	Regulation 2026/1744, recital 7	n/a	Outside the safety function where the purpose is solely efficiency
Credit and eligibility decisions in Colorado	SB 26-189, as reported by law firms	1 Jan 2027	Notice duties for automated decisions
AI-generated calls	FCC 24-84, proposed	Not adopted	Disclosure at the start of each call

Source: the instruments named in each row (Filed); the mapping of candidates to rules is Blue Prysm's reading (Judgement).

Each pilot measures value net of run cost beside a guardrail

The [AI opportunity portfolio model](#), which also exists as a [framework workbench](#), runs a monthly kill gate with six triggers. It stops a pilot when the falsifier fires, when the decision owner leaves or when the data does not exist. It stops or rescopes when the proof runs past twice its plan, re-ranks when the cost to prove doubles, and promotes what works, with an owner and a date. 3GPP's TS 28.105 asks a model learning on a live network for "tolerable degradation thresholds" and "fallback actions", and TS 28.100 gives a human decision "the highest authority in each level". Between them those cover most of it. The failure record adds what they leave open: a guardrail has to be measured beside the value, and some of the rules come from outside the business.

Every pilot therefore measures its value in a unit the business already tracks, net of what the AI costs to run, and beside it measures whatever gets worse if the AI is wrong (Exhibit 5). Bain's advice in September 2026 was to "Measure cost per task on one workflow", as a "cost per resolved customer issue, network incident, proposal generated, or software release". Its scenario puts AI agent and token costs at 20% to 30% of the total, and a fifth to a third is a share worth netting off. Vodafone states its €100 million in [part 2](#) as a net figure.

EXHIBIT 5

Each candidate has a value measure in a unit the business tracks, a guardrail that worsens if the AI is wrong, and a control

CANDIDATE	VALUE, IN	GUARDRAIL BESIDE IT	CONTROL
Energy saving	kWh metered at the site	User throughput and access delay at low load	Matched sites
Agent assist	Cost per resolved contact	Follow-up contact; complaints	Agent groups split at random
Deflection	Contacts resolved without a person, defined in public	Repeat contact; time to reach a person	Customers split at random
RAN optimisation	Spectral efficiency; throughput	Drops and access failures, neighbours included	Matched clusters
AIOps	Repair time on matched incident classes	Critical alarms missed in a shadow run	Shadow run
Fraud and scam	Losses avoided	Legitimate traffic blocked	Labelled holdout
Churn and offers	Churn and margin	Offer cost; complaints	Holdout
Software	Cycle time; defects and rollbacks	Incidents from generated code	Comparable teams
Automated change	Changes completed without a person	Rollback time; local post-change KPIs	Staged rollout

Source: pairing by Blue Prysm, from Bain (10 September 2026), Verizon's customer survey, 3GPP TR 38.864, the FCC report on AT&T and the Optus review (Judgement).

The stop rules

The stop rules are written before the pilot starts, so that stopping is arithmetic rather than argument. Each of the eight comes from an outage, a regulation or a 3GPP control described above (Exhibit 6).

EXHIBIT 6

Eight stop rules, each drawn from an outage, a regulation or a 3GPP control

- 1. Change control.** Stop any tool, AI or otherwise, that can lower a change's risk class or load a change that has not passed review. Rogers' algorithm rated its change Low, and the FCC wrote that loading such a change "should not be possible".
- 2. Degradation.** Roll back when a guardrail crosses the threshold set before the pilot, which is what TS 28.105 asks of models learning live. Roll back too when the post-change check cannot see the local measure, as happened at Optus.
- 3. Care.** Stop a deflection pilot when repeat contacts or complaints rise above the control, whatever the containment rate says.

4. **Consent and disclosure.** Stop an outbound AI voice campaign that lacks documented prior express consent, and stop an EU-facing agent that does not say it is an AI.
5. **Net value.** Stop when value net of model, token and integration cost falls below the bar set before scoring; the portfolio model's bar is three times the cost of proof in a year. Vodafone turned off hundreds of use cases whose value did not justify their cost, and it is better to do that early.
6. **Energy.** Stop a sleep policy when user throughput at low load falls below its threshold, however many kilowatt-hours it is saving. TR 38.864 records a user-throughput loss of up to 62.4% in one source's simulation of one technique.
7. **The model's gate.** Stop when the falsifier fires, the owner leaves or the data is not there. Rescope when the proof runs past twice its plan, and re-rank when its cost doubles.
8. **Dates.** Re-check any credit or eligibility use in Colorado before 1 January 2027, and the classification of any automated-change system before 2 December 2027.

Source: the outages, rules and 3GPP controls cited in Exhibits 1 to 4 and the AI opportunity portfolio model; the rules themselves are Blue Prysm's (Judgement).

Terms used

- **TCPA.** The Telephone Consumer Protection Act, which restricts calls that use an artificial or prerecorded voice.
- **Annex III, safety component.** The AI Act's list of high-risk uses, and the part of a system whose failure endangers people or property.
- **Method of procedure.** The written steps for a network change, which a peer reviews before the change is loaded.
- **Falsifier.** In the portfolio model, the result written down before a pilot starts that will stop it.

Implications

CARRIER STRATEGIST

Put the gate in change control before you choose a model. No AI that touches configuration may lower a risk class or load an unreviewed change, and rollback is automatic, because those are the steps where Rogers, AT&T and Optus failed. Then write each pilot's falsifier and guardrail before it starts.

INVESTOR

Price the precedents as outages rather than as AI fines, because no regulator has fined an operator for its own AI. The dates to hold are 1 January 2027 in Colorado and 2 December 2027 in the EU, with chatbot disclosure in force since August 2026. The question to ask a carrier is which of its AI systems can change the network without a person, and what stops them.

VENDOR

Ship the controls with the model: TS 28.105 degradation thresholds and fallback actions, change records a peer can review, and post-change measures at the level of the cell. A product the customer's change process cannot stop will not get through that process.

Method and limits

HOW THIS WAS BUILT

The rules are taken from the instruments themselves: the FCC's ruling, notice and forfeiture order, the executive order, the Colorado bill pages and the EU regulations in the Official Journal. Law-firm summaries and trackers fill in where the primary record lags behind. The failure record comes from regulators' reports, from the consultant's report the CRTC published on Rogers and, where no report has been published, from the press. The measures and the stop rules are Blue Prism's, built from that record, from 3GPP TS 28.105 and TS 28.100 and from the [portfolio model's](#) kill gate.

WHAT IT DOES NOT SHOW

This is a reading of the public record and not legal advice. Whether a system is high-risk under the AI Act turns on its intended purpose, which only its owner can state. The FCC's proposed disclosure rule, the executive order's deliverables and the FTC's policy statement were all still pending at the trackers' last report. Only law firms report Colorado SB 26-189's signing and effective date; the legislature's own page showed nothing beyond its introduction, on 1 May 2026, at the last check. No FCC report on Verizon's January 2026 outage and no enforcement outcome for AT&T's of February 2024 had been published.

Data as of: rules, regulators' reports and trackers as read 7 Oct 2026 · Method version 1.0.

Sources. FCC 24-17, the 24-84 notice as published in the Federal Register, 24-104 and the Bureau's report on the AT&T outage (22 July 2024); CyberScoop on Lingo Telecom; Xona Partners' assessment of Rogers for the CRTC (4 July 2024); the ACMA on Optus; The Register, The Nightly and Optus's FY26 release; Light Reading on Verizon and Comcast; the Wisconsin Examiner on the Senate vote; Executive Order 14365; the FTC (1 July 2026); Colorado SB25B-004 and SB26-189, with the law firms Seyfarth and McDermott; Regulations (EU) 2024/1689 and 2026/1744 and Directive 2022/2557; NIST AI 100-1; the trackers Governome and The Political Group; Bain (10 September 2026). Read 7 Oct 2026.



WORK WITH BLUE PRYSM

Start with a diagnosis

Blue Prysm is a boutique consultancy working where telecommunications, AI and business strategy meet. We take the work from diagnosis and ideation, through business model design and market testing, to a first commercial offering.

Tell us what you are trying to decide. We reply in writing within two working days, with how we would approach it, what it would take and what it would cost.

PRACTICE

Telecommunications strategy

Spectrum positions, satellite and direct-to-device plans, 6G roadmaps, private networks and the economics of each, priced from filings and auction records.

PRACTICE

AI that pays its way

An AI programme designed around a portfolio of tests, each with its cost to prove, its guardrail and its kill gate written down before it starts.

PRACTICE

Business strategy

Diagnosing a revenue plateau, finding and testing the growth opportunities, and designing the plan to enter them.

How an engagement runs

The first step is a one-day teardown. Each step stands on its own, has a fixed fee agreed before it starts, and you decide whether the next one is worth doing.

STEP 1

One-day teardown

A 90-minute working session on the decision in front of you, a written answer in two working days, and the model behind it, yours to keep.

STEP 2

Two-week assessment

Two working sessions, a 12 to 18 page memo with the working shown so your team can check every number, and a one-page board summary.

STEP 3

Project

Scoped to the engagement, with weekly checkpoints, until a first offering is in the market with your team running it. Every model and source is yours at the close.

The other series

Each series takes one subject and answers five to eight questions about it from the primary record. All are free to read on the site or to download as ebooks.

Seven questions about direct-to-device

Two satellite operators may transmit on American mobile spectrum, and both rent it.

Seven parts

[Read the series →](#)

Eight questions about 6G

6G has a calendar with meeting numbers on it: specifications frozen in December 2028, the first American auction in April 2027 and the conference that settles the spectrum in late 2027.

Eight parts

[Read the series →](#)

Five questions about integrated sensing

A mobile network that senses is a radar, and a weak one beside a purpose-built radar: a metre and a half of range resolution against the metre the standard asks for, and echoes that fade with the fourth power of distance.

Five parts

[Read the series →](#)

Five questions about non-terrestrial networks

Altitude makes the first four design decisions before anyone at a satellite company makes one, so the stratosphere, low orbit, medium orbit and the geostationary arc are four different businesses.

Five parts

[Read the series →](#)

Five questions about private networks and neutral host

The shared band at 3.5 GHz was sold as the home of the private network, and mobile and cable operators paid 86% of the money.

Five parts

[Read the series →](#)

The framework library

132 strategy frameworks, each with a workbench that runs in the browser, and the working files: spectrum valuation, capacity to invest, partnership go/no-go and the AI opportunity portfolio.

Frameworks
· tools

[Open the library →](#)

www.blueprism.com/analysis/ and www.blueprism.com/tools/. © 2026 Blue Prism. Free to use and adapt; attribution with a link to www.blueprism.com is asked for, not required.