Nine industries. One number each.

How AtScale settles the metric every team argues about, so people, BI tools and AI agents get the same answer.

ATSCALE

A metric is a number plus the rules behind it.

Every industry has a metric that two teams calculate in two reasonable ways and get two different answers. An AI agent picks one of them and won't tell you which. AtScale is the semantic layer that defines each metric once, so every person, BI tool and AI agent uses the same definition.

ATSCALE AVA · ASK, VERIFY, ANSWER

Resolves the question

A prompt isn't a specification. Ava works out which metric and which rules a question means before anything is computed, so a vague question gets settled first.

ATSCALE ACE · AI COMPUTATION ENGINE

Computes the metric

ACE calculates the metric from the governed definition, in your warehouse, the same way in every tool, and returns the query it ran alongside the number.

IN THIS GUIDE

01	Retail and consumer	Same-store sales	03
02	Financial services	Net interest margin	05
03	Healthcare and life sciences	Thirty-day readmission rate	07
04	Insurance	Claim count	09
05	Media and entertainment	View count	11
06	Manufacturing	Downtime	13
07	Supply chain and logistics	On-time delivery	15
08	Telecom	Churn	17
09	Adtech	Conversion count	19

AtScale for retail and consumer: one comp number for every merchant, planner and agent.

AtScale is the semantic layer that defines each retail and consumer metric once, so AI assistants, agents and BI tools give merchandising and finance the same number. When someone asks about comps, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do retail teams need one number?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Merchandising and inventory

Sell-through, stock position and SKU performance, consistent across systems and regions.

FEATURED EXAMPLE · SAME-STORE SALES

02

Customer 360

Marketing and merchandising use the same definition of customer value, across CRM, POS, ecommerce and loyalty.

03

Omnichannel

Store, online and marketplace sales compared on one basis, from one model.

04

Demand forecasting

Forecasting and planning models read the same inputs the reports do.

05

Finance

Finance closes the month on the same sales, margin and inventory figures the merchants report.

06

Promotions and markdowns

Promotion lift and margin after markdowns measured the same way by merchants and finance.

How does AtScale keep same-store sales the same for every AI agent?

Comps are up for merchandising and flat for finance. Earnings guidance goes out Thursday, and the agent's already picked one.

Merchandising, finance and store operations each calculate same-store sales differently, and each method is reasonable. An AI agent picks one and won't tell you which, then uses it to decide what to reorder and what to mark down. AtScale defines it once: the tenure threshold, the remodel rule, the closure rule and the 4-5-4 calendar.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
"How did sales grow?"	Same-store sales, stores open a full year, remodels kept in, on the 4-5-4 calendar, against the same weeks last year.
"Which SKUs are losing us money in the Southeast?"	SKUs whose margin after markdowns fell this quarter, in the Southeast as finance draws it.
"Did the promotion work?"	Did basket size rise during the promotion, net of demand pulled forward from the following month?

Jobs AI agents run on same-store sales

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY MONDAY "Post weekly comps to the store managers' channel."	Same-store sales on the 4-5-4 calendar, stores open a full year, remodels kept in, against the same weeks last year.	FULLY AUTOMATED The agent posts on its own. A wrong post is easy to correct.
EVERY NIGHT "Mark down any Southeast SKU whose margin after markdowns fell this quarter."	Margin after markdowns by SKU, in the Southeast as finance draws it.	PERSON APPROVES The agent drafts the markdown list. A merchant approves it before prices change.
END OF EACH PROMOTION "Rebook the promotions that paid for themselves."	Basket-size lift during each promotion, net of demand pulled forward from the next month.	PERSON APPROVES The agent proposes next month's promo calendar for a planner to sign off.

AtScale for financial services: margin you can explain to an examiner.

AtScale is the semantic layer that defines each financial services metric once, so AI assistants, agents and BI tools give treasury and finance the same number. When someone asks about margin, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do a bank's numbers have to agree?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Profitability and margin

Spread measured by customer, product, line of business and region, on one set of transfer-pricing rules.

FEATURED EXAMPLE · NET INTEREST MARGIN

02

Portfolio and research

Position and performance mean the same thing on every desk and in every region.

03

Risk and regulatory

Risk figures that hold up under examination, traceable from source to filed number.

04

Customer 360 and marketing

Prospect, client and counterparty resolved into one view across systems that each hold a piece.

05

Fraud and financial crime

Analysts and scoring models work from one definition, so a case triages consistently.

06

Cost and cloud governance

Acceleration cuts warehouse spend, and teams stop rebuilding each other's cubes.

How does AtScale keep net interest margin traceable for AI agents?

Regulators don't accept “the AI said so.”

A bank has to explain every number it reports to examiners who can fine it. Average earning assets is averaged over the period, and margin can't be rolled up, so a region isn't the sum of its products. AtScale computes each level from its own numerator and denominator, so every number traces back to its calculation.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
“What's our margin by region?”	Net interest margin by region, each one from its own net interest income over its own average earning assets.
“What's our exposure to this counterparty?”	Net exposure to this counterparty across every desk, on today's positions.
“Where did the loss provision move?”	Which segments moved the loss provision this quarter, and by how much, against last quarter's close.

Jobs AI agents run on net interest margin

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY MORNING **“Flag any region whose net interest margin slipped more than 10 basis points.”**	NIM by region, each from its own net interest income over its own average earning assets.	FULLY AUTOMATED The agent sends the alert on its own. Nothing changes until someone reads it.
EVERY HOUR **“Hold new trades with any counterparty over its exposure limit.”**	Net exposure to each counterparty across every desk, on current positions.	PERSON APPROVES The agent puts the trade on hold and pages the risk officer, who releases or blocks it.
AT QUARTER CLOSE **“Draft the loss-provision commentary for the CFO.”**	Which segments moved the loss provision, and by how much, against last quarter's close.	PERSON APPROVES The agent writes the draft. Finance checks every number before it goes out.

SEE IT WITH YOUR METRICS Bring your hardest net interest margin question. **Book a demo →**

ATSCALE

AtScale for healthcare and life sciences: a readmission rate every facility can sign.

AtScale is the semantic layer that defines each healthcare and life sciences metric once, so AI assistants, agents and BI tools give care teams and finance the same number. When someone asks about readmissions, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do care teams and finance need the same figure?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Population health

Cost and outcome tracked per member across payer and provider data.

FEATURED EXAMPLE · THIRTY-DAY READMISSION RATE

02

Clinical and patient analytics

EHR, claims and trial data in one view, with outcome measures that agree across care teams.

03

Revenue cycle and claims

Denials, collections and reimbursement, reconciled on both sides of the contract.

04

Compliance and privacy

HIPAA access rules enforced once, at the metric layer, for every report.

05

Operational efficiency

Operating measures shared across facilities, reachable without a data team in the middle.

06

Cost per case

Cost per case adjusted for case mix, so facilities compare on equal terms.

How does AtScale keep readmission rates consistent for AI agents?

A patient discharged in June comes back in July. Does that count as a readmission, and which month owns it?

Thirty days runs from each discharge, so it never lines up with a calendar month. Planned returns, transfers and deaths all come out, and each exclusion is written policy. AtScale defines those rules once, because mistakes in them end up in regulatory filings.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
"How many readmissions in July?"	Unplanned readmissions within thirty days of a June or July discharge, counted against the discharge month.
"What's our readmission rate?"	Thirty-day readmission rate by facility and service line, with planned returns, transfers and deaths taken out.
"What does a case cost us?"	Cost per case adjusted for case mix, so the facilities with sicker patients aren't penalized.

Jobs AI agents run on thirty-day readmission rate

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY MORNING "Queue a follow-up call for every patient at risk of readmission."	Discharges in the last thirty days, with planned returns, transfers and deaths taken out.	FULLY AUTOMATED The agent adds the calls to the care team's worklist on its own.
EVERY MONTH "Send each facility its readmission rate against target."	Thirty-day readmission rate by facility and service line, counted against the discharge month.	FULLY AUTOMATED The agent sends the report on its own. Each facility sees its own number.
EVERY WEEK "Flag facilities whose cost per case is rising."	Cost per case adjusted for case mix, so facilities with sicker patients aren't flagged for it.	PERSON APPROVES The agent flags the facility. A finance lead decides whether to open a review.

SEE IT WITH YOUR METRICS Bring your hardest thirty-day readmission rate question. Book a demo →

ATSCALE

AtScale for insurance: count a claim once, everywhere.

AtScale is the semantic layer that defines each insurance metric once, so AI assistants, agents and BI tools give claims and actuarial the same number. When someone asks about claims, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do claims, pricing and reserving need one count?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Claims optimization

Triage and cycle time from one claim record, with fraud flags that agree across teams.

FEATURED EXAMPLE · CLAIM COUNT

02

Per-member cost

Cost per member per month means the same thing to payer and provider teams.

03

Regulatory reporting

Figures ready for state and federal filings, traceable back to source.

04

Customer 360 and retention

Agents, call center and digital resolved into one policyholder, with retention measured once.

05

Actuarial and underwriting

Pricing and rating models read the same risk data the product team reports on.

06

Reserving and catastrophe response

Reserves and catastrophe exposure updated from the same claim data the actuaries use.

How does AtScale keep claim counts the same for AI agents?

A claim reopened in August gets counted as new in September, and your loss ratio just improved.

Adjusters and actuaries set the rules for how claims are counted, and those rules often live only in individual analysts' queries. Accident year and calendar year answer differently, and frequency and loss ratio inherit whichever rule was used. AtScale defines the count once, so an agent working the claims queue applies the same rules to every file.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
“How many claims did we open in 2025?”	Claims on accident year 2025, counting a reopened file once.
“Is frequency up?”	Claim frequency this year against last, on the same reopen rule in both years.
“How big was the catastrophe?”	Claims tied to that event, including the ones closed without payment.

Jobs AI agents run on claim count

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY MORNING “Assign new claims to adjusters by region.”	New claims on accident year, with a reopened file counted once.	FULLY AUTOMATED The agent routes the queue on its own. A claim is easy to reassign.
EVERY MONTH “Alert pricing when claim frequency rises in any line.”	Claim frequency this year against last, on the same reopen rule in both years.	FULLY AUTOMATED The agent sends the alert on its own. Pricing decides what to do with it.
EVERY MORNING “Update the reserve for new claims tied to the Gulf storm.”	Claims tied to that event, reopened files counted once, closed-without-payment files included.	PERSON APPROVES The agent proposes the reserve change. An actuary approves it before it posts.

SEE IT WITH YOUR METRICS Bring your hardest claim count question. Book a demo →

ATSCALE

AtScale for media and entertainment: one view count for ad sales, programming and royalties.

AtScale is the semantic layer that defines each media and entertainment metric once, so AI assistants, agents and BI tools give ad sales and finance the same number. When someone asks about views, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where does every screen need the same audience number?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Audience and viewership

Linear, streaming and social reconciled into one audience figure.

FEATURED EXAMPLE · VIEW COUNT

02

Content performance

Titles compared across genres, platforms and release windows on equal terms.

03

Subscribers and churn

Billing, app and CRM agree on what a subscriber is, so churn means one thing.

04

Rights and royalties

Licensing and residual figures that finance, legal and content all read alike.

05

Ad sales and monetization

Inventory and yield on the same footing whether the sale is direct or programmatic.

06

Chat and agents

Chat and agents pull answers from the metric layer, so nobody writes SQL against raw tables.

How does AtScale keep view counts the same for AI agents?

One password can cover four households, and Wall Street still expects a single subscriber count.

A view is whatever completion rate the business set, and moving it reorders every title. A billing rule decides whether a shared account is one subscriber or four. AtScale defines both once, because ad rates, commissioning calls and royalty checks all move with them.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
"Which title got the most views?"	Views at the thirty-percent completion threshold, streaming and linear together, last calendar month.
"How many subscribers do we have?"	Paying accounts, counted once per account rather than once per profile.
"What do we owe on this title?"	Royalties owed on this title, with streaming and linear counted on one definition.

Jobs AI agents run on view count

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY MONDAY **"Promote the top ten titles on the home screen."**	Views at the thirty-percent completion threshold, streaming and linear together.	FULLY AUTOMATED The agent updates the home screen on its own. It's easy to roll back.
EVERY DAY **"Send a win-back offer to accounts that stopped paying."**	Paying accounts, counted once per account rather than once per profile.	FULLY AUTOMATED The agent sends offers on its own, inside a budget someone set.
MONTH END **"Draft royalty statements for every rights holder."**	Views per title, with streaming and linear on one definition.	PERSON APPROVES The agent drafts the statements. Rights and finance approve them before payment.

SEE IT WITH YOUR METRICS Bring your hardest view count question. **Book a demo →**

AtScale for manufacturing: rank every plant on the same clock.

AtScale is the semantic layer that defines each manufacturing metric once, so AI assistants, agents and BI tools give plant managers and headquarters the same number. When someone asks about uptime, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do plants need to measure the same way?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Production efficiency

Bottlenecks and lost hours read out of MES the same way in every plant.

FEATURED EXAMPLE · DOWNTIME

02

Supply chain and inventory

Live inventory across ERP and warehouse systems, with the same measures plant to plant.

03

Quality control

Quality standardized across plants and product lines, and ready when an auditor asks.

04

Predictive maintenance

Sensor data joined to ERP, with one equipment-health definition feeding the models.

05

Aftermarket and warranty

Distributors, dealers and service in one view, with warranty comparable by region.

06

Chat and agents

Plant managers and agents ask in plain language and get the same numbers the reports show.

How does AtScale keep downtime the same across plants for AI agents?

Plant 4 has the best uptime because it counts changeovers as running, and you just moved the capex there.

Each plant decided long ago what counts as downtime, and each call made sense locally. The MES, the ERP and the sensors all keep different clocks, and four plants don't average into a region. AtScale defines downtime once, because headquarters' ranking decides where capital goes.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
“Which plant has the best uptime?”	Availability by plant, changeovers counted as downtime everywhere, on the MES clock.
“Which lines lost the most hours?”	Unplanned downtime hours by line last month, split by cause code.
“What did stops cost us?”	The cost of unplanned downtime across the region this quarter, at each line's own rate.

Jobs AI agents run on downtime

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY SHIFT **“Rank the plants on uptime and post it to the ops channel.”**	Availability by plant, changeovers counted as downtime everywhere, on the MES clock.	FULLY AUTOMATED The agent posts the ranking on its own.
EVERY NIGHT **“Open a maintenance ticket for any line that lost more than four hours.”**	Unplanned downtime hours by line, split by cause code, planned maintenance left out.	FULLY AUTOMATED The agent opens the ticket on its own. A ticket is cheap to close.
EVERY WEEK **“Move spare parts to the lines where stops cost the most.”**	The cost of unplanned downtime by line, at each line's own rate.	PERSON APPROVES The agent proposes the transfer. A plant manager approves it before parts move.

AtScale for supply chain and logistics: on time means the date you promised.

AtScale is the semantic layer that defines each supply chain and logistics metric once, so AI assistants, agents and BI tools give transportation and customer service the same number. When someone asks about on-time delivery, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do planners, carriers and customers need one answer?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Service performance

On time by carrier, distribution center, customer and product, all from one rule.

FEATURED EXAMPLE · ON-TIME DELIVERY

02

Demand and supply planning

Forecasts and the reports that grade them built from the same history and the same calendar.

03

Inventory optimization

Stock on hand, days of cover and stockout risk, consistent across every distribution center.

04

Supplier performance

Lead time, fill rate and defect scorecards that procurement and finance both accept.

05

Transportation and network

The service and cost impact of a reroute, answered before anyone moves a load.

06

Cost to serve

Freight, handling and expedite allocated to customer and product by one rule.

How does AtScale keep on-time delivery the same for AI agents?

The carrier logged the shipment as on time, and the customer got it late.

On time depends on whose date counts and whether a partial shipment counts, and months don't average into a quarter. Roll it up like revenue and the total looks fine but it's wrong. AtScale defines it once, so a routing agent never moves freight on the carrier's number.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
"What's our on-time rate?"	Deliveries meeting the date we promised, partials counted late, computed for the quarter rather than averaged from the months.
"Which carriers are hurting us?"	On-time rate by carrier, each drawn from its own shipments over the same period.
"What happens if I reroute through Memphis?"	The service and cost effect on the East Coast if that volume moves to Memphis.

Jobs AI agents run on on-time delivery

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY MORNING "Send each customer its on-time rate for the quarter."	Deliveries meeting the date we promised, partials counted late, computed for the quarter.	FULLY AUTOMATED The agent sends the scorecard on its own.
EVERY NIGHT "Move volume off any carrier that drops below 92% on time."	On-time rate by carrier against the promised date, each from its own shipments.	PERSON APPROVES The agent drafts the reroute. A transportation manager approves it before loads move.
ON REQUEST "Model rerouting East Coast volume through Memphis."	The service and cost effect of moving that volume, on the same on-time definition as the carrier scores.	PERSON APPROVES The agent returns the scenario. A person decides whether to reroute.

AtScale for telecom: churn that matches the bill.

AtScale is the semantic layer that defines each telecom metric once, so AI assistants, agents and BI tools give marketing and finance the same number. When someone asks about churn, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where do billing, network and marketing need to count the same subscriber?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Subscribers and churn

One subscriber definition across billing, CRM and provisioning, with suspends and downgrades ruled on once.

FEATURED EXAMPLE · CHURN

02

Network performance

Counter data from 3G, 4G and 5G hardware, comparable across vendors.

03

Customer experience

Network events tied to the problems customers report, so support and engineering trace one root cause.

04

Self-service analytics

Departments reach the metrics directly, without a ticket to the central data team.

05

Cost optimization

Automated aggregates cut query cost on every platform the queries run on.

06

Governance and security

Role-based access at the metric layer, so regional and vendor systems stop carrying their own copy of the logic.

How does AtScale keep churn the same for AI agents?

A suspended line still gets a bill, but your churn number says it left.

Each team rules suspends, downgrades and plan migrations in or out a little differently, and the subscriber base moves every month. Neither shows up until finance checks the numbers against billing. AtScale defines churn once, against the average base, so every agent reports the figure finance will sign.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
“What was churn last quarter?”	Churn computed across the quarter against the average subscriber base, suspends left out.
“Which plans are churning?”	Churn by plan and tenure cohort, each cohort against its own base.
“What's the margin on this plan?”	Plan margin with the handset subsidy loaded in.

Jobs AI agents run on churn

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
AT QUARTER CLOSE “Report churn to finance for the board deck.”	Churn across the quarter against the average subscriber base, suspends left out.	PERSON APPROVES The agent drafts the slide. Finance signs off before it reaches the board.
EVERY DAY “Send a retention offer to any plan cohort where churn is rising.”	Churn by plan and tenure cohort, each cohort against its own base.	FULLY AUTOMATED The agent sends offers on its own, capped per customer.
EVERY WEEK “Flag any plan whose margin went negative.”	Plan margin with the handset subsidy loaded in.	PERSON APPROVES The agent flags the plan. Pricing decides whether to change the rate card.

AtScale for adtech: one conversion count, whoever sends the invoice.

AtScale is the semantic layer that defines each adtech metric once, so AI assistants, agents and BI tools give media buyers and finance the same number. When someone asks about conversions, **AtScale Ava** (Ask, Verify, Answer) checks which version they mean. **AtScale ACE**, the AI Computation Engine, then calculates it the same way in every tool.

USE CASES

Where does every platform need to count the same way?

AtScale settles each definition once, and every person, tool and agent uses it. Without that, two teams can calculate the same metric two reasonable ways and get two different answers.

01

Attribution and measurement

One attribution model for marketing, sales and finance.

FEATURED EXAMPLE · CONVERSION COUNT

02

Campaign performance

Impressions, clicks and conversions defined once across DSPs, SSPs and publishers.

03

Audience segmentation

Segments defined by data science and media together, activated without being rebuilt.

04

Programmatic and bid optimization

Bid, win rate and yield lined up across exchanges, and fed live to pacing models.

05

Privacy and identity

Access-controlled definitions, with identity resolution that holds from platform to platform.

06

Invoice reconciliation

Platform invoices checked against your own conversion and impression counts before you pay.

How does AtScale keep conversion counts the same for AI agents?

Three platforms took credit for one sale, and all three invoiced you.

Each platform picks its own action, its own attribution model and its own window, then bills on its own count. At quarter end you're left with three numbers to reconcile. AtScale defines the conversion once and applies it to every platform's raw events.

What people ask, and what they mean

People type short questions. Before AtScale ACE computes anything, AtScale Ava works out which metric and which rules each question means.

WHAT A PERSON TYPES	WHAT THEY MEANT
"How many conversions did that campaign drive?"	Purchases, last non-click credit, a seven-day window, the same rule on every platform.
"Why is the DSP above our number?"	The gap between the DSP's count and ours, once both run our action, our credit rule and our window.
"Which segments work?"	Segments that perform across platforms once conversions are deduplicated.

Jobs AI agents run on conversion count

When an agent asks, Ava applies the definition the business already confirmed. Fully automated means a low-stakes action the agent takes on its own. Person approves means someone signs off first.

WHAT THE AGENT IS TOLD	WHAT ACE CALCULATES	WHO ACTS
EVERY HOUR "Pause any campaign above its target cost per conversion."	Purchases on last non-click credit and a seven-day window, the same rule on every platform.	FULLY AUTOMATED The agent pauses on its own. A paused campaign restarts in a click.
MONTH END "Dispute the DSP invoice when its count runs above ours."	The gap between the DSP's count and ours, once both run our action, our credit rule and our window.	PERSON APPROVES The agent drafts the dispute. Finance reviews the dispute and sends it.
EVERY HOUR "Shift budget to the segments converting best."	Segment conversions deduplicated across platforms, so the agent doesn't pay three times for one sale.	FULLY AUTOMATED The agent shifts budget on its own, inside limits a person set.

SEE IT WITH YOUR METRICS Bring your hardest conversion count question. Book a demo →

ATSCALE

Bring your hardest metric.

For over a decade, Fortune 1000 companies have relied on the metrics AtScale computes. Bring yours, and we'll work through them with you for your AI agents.

STEP 01

Bring your questions

The ones your team asks every week and argues over every quarter.

STEP 02

We run the benchmark

Your AI answers them twice, once with AtScale under it and once without, and you read both.

STEP 03

You keep the model

The semantic model we build for the test is yours, whatever you decide afterward.

Book a demo atscale.com/demo

ATSCALE