

Human-Verified vs. AI-Verified Contacts

Accuracy benchmarks for 2026, and why a single accuracy percentage
hides the decision that matters

21%

of automatically verified records are deliverable and still wrong,
reaching a real mailbox belonging to the wrong person

MODELED, AGGREGATE VERIFICATION DATA

2026 EDITION

Research by Database Providers

July 2026 · thedatabaseproviders.com

TWO METHODS ANSWERING TWO DIFFERENT QUESTIONS



ABOUT THIS REPORT

This report benchmarks automated verification against human-assisted verification for B2B contact data. It separates two properties that are routinely reported as one number, models the cost of each method once downstream handling is counted, and sets out where each is the correct choice.

The central argument is that a single accuracy percentage is not a usable comparison, because the two methods fail at different things. Automated verification establishes whether an address will accept mail. Human verification establishes whether the person named still holds the stated role at the stated organisation. A record can pass the first test completely and fail the second, and on the modeling here that describes 21% of an automatically verified file.

The report covers the accuracy decomposition, the four failure modes automation is structurally unable to detect, a cost model with stated inputs, a tiering approach based on downstream handling intensity rather than record importance, and what vendors actually mean when they describe verification as AI-driven.

PUBLISHED July 2026

EDITION 2026 Edition

RESEARCH BY Database Providers

AUDIENCE CMOs, demand generation leaders, RevOps and marketing operations teams, data procurement buyers

REPORT CODE DP-WP-2026-07

CONTENTS

Foreword

Executive summary

Key findings at a glance

01 Two methods answering two different questions

02 Where automation is structurally blind

03 The single accuracy number hides the decision

04 Tier by handling intensity, not by account importance

05 What vendors mean by AI-verified

The Verification Depth Model

Methodology

About Database Providers

FOREWORD

The question we are asked most often about verification is which method is more accurate, and it is the wrong question in an interesting way. Not because the answer is complicated, but because accuracy is not one property. A contact record makes two separate claims: that this address will reach a mailbox, and that the mailbox belongs to a person holding this role at this company. The methods available to verify those two claims are different, cost different amounts, and are good at different things.

Automation is excellent at the first claim and structurally incapable of the second. An automated stack can open a conversation with a receiving mail server and establish that a mailbox exists. It cannot establish that the person who owns that mailbox still works there, still holds the title recorded against them, or is still the right person to contact. Those facts are not exposed by any protocol. They exist in the world, not in the mail system.

This creates a category of failure that is worse than a bounce, because it is silent. A record that is deliverable and misattributed passes every automated check, arrives successfully, counts as delivered, and reaches somebody for whom the message makes no sense. Nothing in any report distinguishes it from a good record.

We should be transparent about our position. Database Providers uses human confirmation as part of its verification process, so a report concluding that human verification adds value is self-serving on its face. Two things in what follows push against that. The cost model shows automation winning decisively on price per record, by roughly seven to one. And Chapter 4 concludes that most of a typical file does not need human verification at all, and identifies the specific portion that does. If the honest answer were that everything should be human-verified, that would be a better commercial outcome for us and a worse report.

Research Team, Database Providers

Fremont, California

EXECUTIVE SUMMARY

Verification methods should be compared on two axes rather than one: whether the address is reachable, and whether the record is true. Automated stacks perform close to human-assisted ones on the first and substantially worse on the second, and the gap only becomes expensive downstream.

FINDING 01**Automated verification reaches 92% deliverability but only 71% fully usable records.**

Human-assisted verification reaches 97% deliverability and 94% fully usable. The deliverability gap is five percentage points; the usable-record gap is twenty-three. Almost all of the difference sits in attribution rather than reachability.

FINDING 02**21% of automatically verified records are deliverable and still wrong.**

These records reach a real, monitored mailbox belonging to someone who no longer holds the stated role, or to a shared alias, or to a catch-all domain that accepted a generated address. They generate no bounce and appear in reporting as successful delivery.

FINDING 03**Automation is roughly seven times cheaper per record and 5.6 times more expensive per usable contact.**

At a modeled USD 0.004 against USD 0.028 per record, automation wins decisively on price. Once sales handling of misattributed records is counted at USD 3.10 each, total cost per usable contact is USD 1.27 for automation against USD 0.23 for human-assisted verification.

FINDING 04**The break-even sits at USD 0.07 of handling cost per bad record, or about six seconds of attention.**

Below that, automated verification is the cheaper choice. Above it, human-assisted verification is. Any record that a person opens, reads, and acts on has already crossed the threshold, which is why the answer depends on how the data will be used.

FINDING 05**Matching verification depth to downstream handling cuts modeled cost per usable contact to USD 0.07.**

Records routed to sales development justify human verification; records routed to low-touch nurture do not, because nobody spends time on the failures. Tiering by handling intensity rather than by account importance produces the lowest modeled cost of the three approaches examined.

KEY FINDINGS AT A GLANCE

Six numbers that summarise the research findings.

21%

DELIVERABLE BUT MISATTRIBUTED

Share of automatically verified records that reach a mailbox and are still wrong. See Chapter 2.

23pp

USABLE-RECORD GAP

94% human-assisted against 71% automated, almost all of it attribution. See Chapter 1.

5.6x

COST PER USABLE CONTACT

Automated against human-assisted, once downstream handling is counted. See Chapter 3.

\$0.07

THE BREAK-EVEN POINT

Handling cost per bad record at which the two methods are equal, about six seconds. See Chapter 3.

4

BLIND SPOTS IN AUTOMATION

Failure modes no protocol-level check can detect, however sophisticated. See Chapter 2.

\$0.07

PER USABLE, WHEN TIERED

Modeled cost when verification depth is matched to downstream handling. See Chapter 4.

01

Two methods answering two different questions

Reachability and attribution are separate properties, and only one of them is visible to a protocol

HUMAN-ASSISTED 94%



AUTOMATED ONLY 71%



CHAPTER 01

Two methods answering two different questions

A contact record asserts two things simultaneously, and they can be true or false independently of each other. The first assertion is about the mail system: this address exists, accepts mail, and is monitored. The second is about the world: the person named holds this role at this organisation and is therefore a sensible person to contact about this subject.

Automated verification interrogates the first assertion directly and well. Syntax checks eliminate malformed addresses. Domain and MX checks confirm the domain exists and is configured to receive mail. SMTP validation opens a conversation with the receiving server to establish whether the specific mailbox exists. Catch-all detection identifies domains that accept everything regardless of local part. Each of these is mechanical, fast, and cheap, and together they establish reachability to a high standard.

None of them touches the second assertion, and no refinement of the method can. Whether a person still works somewhere is not a property exposed by SMTP, by DNS, or by any protocol a verification stack can query. It is a fact about employment, and establishing it requires looking at the world rather than at the mail system.

Modeled across the two methods, the consequence is a gap that is small on the first axis and large on the second. Automated stacks reach 92% deliverability against 97% for human-assisted verification, a five-point difference. On attribution accuracy among deliverable records, automation reaches 77% against 97%, a twenty-point difference.

The methods are close on reachability and far apart on whether the record is actually true.



Source: Database Providers, modeled verification accuracy by method, 2026. Attribution measured among deliverable records.

Why a single accuracy figure misleads

Vendors quote one accuracy number, and buyers compare those numbers. The practice survives because it is convenient and because there is no agreed definition of what the number measures, which means a vendor quoting 95% accuracy may be describing deliverability, usable records, or something in between.

The ambiguity is not always deliberate, but it is consistently favourable to automated methods, because deliverability is both the easier property to achieve and the easier one to measure. A vendor with a strong automated stack and no human layer can honestly report a high accuracy figure that says nothing about whether the titles in the file are current.

The practical defence is a question rather than a benchmark: ask what the quoted accuracy figure measures, and specifically whether it includes the claim that the person holds the stated role at the stated organisation. A vendor who separates the two numbers is telling you they measure both. A vendor who cannot separate them is telling

you they measure one.

Whether a person still works somewhere is not exposed by any protocol. It is a fact about employment, not about the mail system.

Combining into a usable-record rate

The figure that matters operationally is neither of the two axes alone but their combination: the share of records that are both reachable and true. Applying the modeled attribution rates to the modeled deliverability rates gives 71% fully usable records for automated verification and 94% for human-assisted verification.

The remaining 29% of an automated file divides unevenly. Eight percent is undeliverable and will bounce, which is visible, actionable, and gets suppressed. The other 21% is deliverable and misattributed, which is invisible and does not.

That 21% is the subject of the next chapter, because it is where the entire cost difference between the two methods originates. If misattributed records were as visible as bounces, the choice between methods would be a straightforward price comparison and automation would win it.

WHAT THIS MEANS FOR LEADERS

- Ask every vendor to state deliverability and attribution accuracy as two separate numbers.
- Treat a single quoted accuracy figure as describing reachability unless proven otherwise.
- Measure your own usable-record rate as reachable and correctly attributed, not as delivered.
- Expect automation to close the reachability gap and not the attribution gap, regardless of sophistication.

02

Where automation is structurally blind

Four failure modes that pass every protocol-level check, however sophisticated the stack

HUMAN-ASSISTED 94%



AUTOMATED ONLY 71%



CHAPTER 02

Where automation is structurally blind


The word structurally is doing real work in this chapter's title. The argument is not that current automated verification is insufficiently advanced and will improve. It is that four specific failure modes produce records which are genuinely, verifiably deliverable, and no amount of protocol-level checking can distinguish them from good records, because at the protocol level they are indistinguishable.

The first and largest is the departed employee whose mailbox remains active. Organisations routinely keep a mailbox live after someone leaves, forwarding to a successor or a shared inbox for a transition period that is often measured in months. The address exists, accepts mail, and is monitored. Every automated check passes. The record is wrong.

The second is the person who is still employed but no longer in the role. Promotions and internal moves change what somebody is responsible for without changing their address. The record is correct about the person and wrong about the reason for contacting them, which produces outreach that is well-addressed and irrelevant.

The third is the shared or role-based alias, where the address resolves to a function rather than a person. The fourth is the generated address accepted by a catch-all domain: where a scrape or an enrichment tool constructs an address from an observed company format, a catch-all domain will accept it regardless of whether the mailbox exists. Catch-all detection can flag that the domain accepts everything; it cannot tell you whether this particular address is real.

Only eight percent of an automatically verified file fails visibly. Twenty-one percent fails silently.



Source: Database Providers, modeled composition of an automatically verified B2B file, 2026.

Why the silent failures cost more than the loud ones

An undeliverable record announces itself. It bounces, enters a suppression list, and never consumes attention again. It damages sender reputation, which the 2026 report on bounce economics prices in detail, but it is a known quantity that can be measured and managed.

A misattributed record does none of that. It is delivered, so it counts as a success. It is opened occasionally, because a real person receives it. It consumes sales development time, because someone follows up on a delivered message with no reply. It sits in the CRM as an active contact at a target account, giving false confidence that the account is covered. And it never generates an error that would prompt anyone to check.

The cost is therefore concentrated in exactly the place where it is hardest to see: sales time spent on records that a report will never flag. That is why the cost model in the next chapter charges misattributed records the same handling cost as undeliverable ones. From the perspective of a person working a queue, a contact who left eight

months ago and a contact who never existed consume the same effort before being written off.

A contact who left eight months ago and a contact who never existed consume the same effort before being written off.

What human confirmation actually does

Human confirmation is a less impressive process than the name suggests, and describing it plainly is useful because it clarifies both its value and its limits. A researcher establishes, from sources outside the mail system, that the named person currently holds the stated role at the stated organisation. That is the whole of it.

Its value is that it addresses the one question automation cannot reach, and it addresses all four failure modes at once: a departed employee is identified as departed, a moved employee as moved, an alias as an alias, and a generated address as unattributable to a real named person.

Its limits are equally clear. It is slow relative to automation, taking days rather than minutes on a large file, which matters for the signal-driven motions described in the 2026 report on intent and contact data. It is roughly seven times more expensive per record. And it is a point-in-time check that begins decaying immediately, at the rate the 2026 data decay report models: roughly 3.05% per month, and faster for the director-level records where attribution matters most.

That last point is the one most often missed in vendor comparisons. Human verification does not produce permanently correct data. It produces data that was correct on a stated date, which is why the date matters as much as the method and why cadence is a separate decision from depth.

WHAT THIS MEANS FOR LEADERS

- Treat deliverable-but-misattributed records as a cost category and estimate their share explicitly.
- Charge misattributed records the same handling cost as bounces, since they consume the same sales time.
- Ask when human confirmation was last performed, not only whether it was performed.
- Do not expect automation to solve attribution; the limit is structural rather than technological.

03

The single accuracy number hides the decision

A cost model with stated inputs, and the six-second threshold that determines the answer

HUMAN-ASSISTED 94%



AUTOMATED ONLY 71%



CHAPTER 03


The single accuracy number hides the decision

Comparing the methods on price produces one answer and comparing them on outcome produces the opposite one. Both comparisons are valid; they simply measure different things, and which one is correct depends on what happens to the records afterwards.

On price, automation wins decisively. At a modeled USD 0.004 per record against USD 0.028 for human-assisted verification, automation is seven times cheaper. For a file of 100,000 records that is USD 400 against USD 2,800, and no procurement process would find that a difficult decision.

On total cost the ranking inverts. Charging each unusable record USD 3.10 in sales development handling, the same figure used throughout this series, a thousand automatically verified records cost USD 4 to buy and USD 899 to work through, for 710 usable contacts. A thousand human-verified records cost USD 28 to buy and USD 186 to work through, for 940 usable contacts. That is USD 1.27 per usable contact against USD 0.23, a factor of 5.6.

For records that a person will actually work, the cheaper file costs four times as much in total.



Source: Database Providers, modeled total cost per 1,000 records including handling at USD 3.10 per unusable record, 2026. Tiering is set out in Chapter 4.

The six-second threshold

A reader who disputes the USD 3.10 handling figure can solve the model for the point where the two methods cost the same. That break-even sits at approximately USD 0.07 of handling cost per unusable record, which is about six seconds of sales development time at a loaded rate of USD 42 per hour.

Six seconds is a useful number because it is so small. Any record that a person opens, reads, forms a judgement about, and dispositions has consumed far more than six seconds. Any record that is loaded into a sequence, sends three messages, and produces no reply has consumed more than six seconds of somebody's attention across the follow-up. In practice, any record touched by a human at all sits above the threshold.

This is what makes the decision tractable. The question is not which verification method is better in the abstract. It is whether a person will spend more than six seconds on each record, and that is a question about the downstream process rather than about the data.

Where the answer is no, automation is genuinely and substantially the better choice, and the price advantage is real rather than an accounting illusion. Records used for audience matching in advertising platforms, for low-touch nurture, or for aggregate analysis never consume human attention on a per-record basis. For those uses the misattribution rate is close to costless, and paying seven times more per record to fix it would be waste.

The question is not which method is better. It is whether a person will spend more than six seconds on each record.

What the model does not capture

Three effects sit outside the model and all of them push in the same direction, which is worth stating so that the figures are read as conservative rather than as complete.

The first is deliverability damage. The 8% of an automated file that is undeliverable produces bounces, and bounces degrade sender reputation for every subsequent campaign from the domain. The 2026 report on bounce economics models this at roughly five suppressed legitimate messages for every message lost to a bounce. None of that appears in the handling cost above.

The second is pipeline coverage. A misattributed record at a target account makes the account look covered when it is not, which distorts account-based programmes in the way the 2026 report on intent and contact data describes. The cost of a signal arriving at an account with no reachable committee member is not a handling cost and is not counted here.

The third is opportunity cost on the message itself. A well-crafted message that reaches the wrong person is not merely wasted; in some cases it is actively unhelpful, because the recipient forms an impression of the sender. We have not attempted to quantify this and would be sceptical of any model that did, but it is not zero.


WHAT THIS MEANS FOR LEADERS

- Compare methods on cost per usable contact, not on price per record.
- Establish whether records will consume more than six seconds of human attention before choosing a method.
- Use automation without hesitation for audience matching, low-touch nurture, and aggregate analysis.
- Treat the modeled figures as conservative, since deliverability and coverage effects sit outside them.

04

Tier by handling intensity, not by account importance

Most of a file does not need human verification, and the portion that does is identifiable in advance



CHAPTER 04

Tier by handling intensity, not by account importance


The obvious way to tier verification is by how much an account matters: human-verify the target accounts and the enterprise segment, automate the rest. It is intuitive, it is what most teams do when they tier at all, and the model says it is the wrong axis.

The right axis is how much human time each record will consume downstream, because that is the only variable that determines whether the extra verification cost is repaid. A record at a small account that an SDR will personally research and call is expensive to get wrong. A record at a strategic enterprise account that will only ever be used for audience matching in an advertising platform is not.

The two axes correlate, which is why tiering by importance produces reasonable results by accident. But they come apart at the edges, and the edges are where the money is. The exception cases are exactly the ones that intuition handles badly.

Modeled across a typical B2B file, the share of records that will be individually worked by a person is smaller than most teams assume. Roughly a fifth goes to sales development for individual handling. Two fifths goes into automated sequences where a person reads only the replies. The remainder is used for nurture, audience matching, or analysis, or is held and never used at all.

About a fifth of a typical file is individually worked by a person, and only that fifth clears the six-second threshold decisively.



Source: Database Providers, modeled allocation of a typical B2B contact file by downstream handling intensity, 2026.

The tiered model

Applying human verification to the quarter of a file that will be individually worked, and automation to the rest, produces the lowest modeled cost of the three approaches. Per thousand records: the human-verified quarter costs USD 7 to verify and generates USD 47 of handling on its 6% failure rate, for 235 usable contacts. The automated remainder costs USD 3 to verify and generates negligible per-record handling cost because nobody works the failures individually, for 533 usable contacts.

Total modeled cost is USD 57 per thousand records for 768 usable contacts, or USD 0.07 per usable contact. Against USD 1.27 for automating everything and USD 0.23 for human-verifying everything, tiering by handling intensity is roughly three times better than the better of the two single-method approaches.

The result depends on an assumption worth making explicit: that the automated portion genuinely does not consume per-record human time. If low-touch records leak into a sales queue, which happens routinely when a nurture list is repurposed for outbound under pressure, the handling cost reappears and the advantage

disappears. The tiering only holds if the routing holds.

This is the practical reason to tier by intended use rather than by account value. Intended use can be recorded as a field on the record and enforced at the point of routing. Account importance is a judgement that changes, and a record tiered by importance carries no instruction about how it may be used.

The tiering only holds if the routing holds. A nurture list repurposed for outbound brings the handling cost back with it.

Implementing the split

Three mechanics make this work in practice. Record intended use as a field at acquisition, so that verification depth is a consequence of a decision that has already been made rather than a separate decision someone has to remember. Verify to the depth that field implies. And enforce the field at routing, so that a record acquired for nurture cannot be loaded into a sales sequence without an explicit upgrade step that triggers verification.

The upgrade path matters as much as the initial split, because the common and legitimate case is a record that starts in nurture, engages, and becomes worth working individually. That transition should trigger human verification of that record at that moment, which is both cheaper than verifying the whole nurture pool and better timed, since the verification happens immediately before the record is used.

Cadence then applies on top of depth. A human-verified record is correct on the date it was checked and decays from there, so the sales development tier needs re-verification on a schedule rather than a one-time check. Quarterly is the cadence the 2026 data decay report supports for most segments, with shorter cycles for volatile verticals and mid-seniority roles.


WHAT THIS MEANS FOR LEADERS

- Tier verification by downstream handling intensity, not by account importance.
- Record intended use as a field at acquisition and enforce it at routing.
- Trigger human verification at the moment a nurture record is promoted to individual working.
- Apply a re-verification cadence to the sales development tier, since human checks decay like any other.

05

What vendors mean by AI-verified

Four different processes share a label, and only some of them address attribution



CHAPTER 05

What vendors mean by AI-verified


AI-verified has become a marketing category rather than a technical description, and it now covers at least four distinct processes with materially different reliability. Distinguishing them is the practical skill a buyer needs, because the label alone carries almost no information.

The weakest is pattern inference: constructing an address from an observed company format and a name, then treating it as verified because it is plausible. This is not verification at all, and it produces exactly the generated addresses that catch-all domains accept. The presence of pattern inference in a vendor's process is the single most useful thing to establish.

The second is machine-learning confidence scoring, which uses historical signals to predict whether an address is likely to be valid. This is genuinely useful and genuinely different from checking: a confidence score is a probability, not an observation, and a file described as verified on the basis of scores contains no records that were actually checked.

The third is model-driven web research, where a system reads public sources to establish employment and role. This addresses attribution directly and is the most promising of the four, with an important caveat: language models produce fluent, plausible output when sources are thin, and a fabricated but reasonable-sounding job title is harder to detect than a malformed address. The safeguard is whether the vendor retains the source for each claim. The fourth is agentic checking, which combines model-driven research with protocol-level validation and is closest to a genuine substitute for human confirmation, subject to the same sourcing caveat.

Most vendors will supply a sample; far fewer can state what was checked, when, or how much of the file was checked at all.



Source: Database Providers, modeled share of B2B data vendors able to answer each diligence question concretely, 2026.

Five questions that separate the categories

The questions that produce useful answers are short and specific, and the pattern of responses is more diagnostic than any individual answer. Ask what share of this file was checked against a live mail server rather than scored. Ask what share had employment and role confirmed from a source outside the mail system, and whether that source is retained. Ask the date each layer was last run against this file. Ask what share of the addresses were constructed from a pattern rather than observed. And ask for deliverability and attribution accuracy as two separate numbers.

Vendors running a genuine multi-layer process answer these quickly and concretely, because the information exists as a by-product of doing the work. Vendors who cannot tend to redirect to a headline accuracy figure, to certifications that do not address method, or to a description of their technology rather than of what was done to

this file.

The last question is the most discriminating and the easiest to ask. A vendor who reports attribution and deliverability separately has, by definition, measured both. On the modeled figures, few can.

As with the other reports in this series, these questions should be put to Database Providers on the same terms as to any other supplier, and an inability to answer them concretely should be treated as disqualifying regardless of who is being asked.

A fabricated but reasonable-sounding job title is considerably harder to detect than a malformed address.

Where this is heading

The honest forecast is that the attribution gap will narrow without closing, and that the mechanism will be model-driven research with retained sources rather than any advance in protocol-level checking. The limit identified in Chapter 2 is structural for mail-system checks and is not structural for systems that read the world; a process that can read public sources is doing the same kind of work a human researcher does.

What will determine whether that closes the gap in practice is verifiability rather than capability. A human researcher can be asked what they checked. An automated research process is only as auditable as the sources it retains, and a system that returns a confident answer without a citation has moved the attribution problem rather than solved it.

For a buyer today, this suggests a specific question to carry into vendor conversations over the next few years: not whether the vendor uses AI, which will shortly be universal and uninformative, but whether each attribution claim carries a retained source that can be inspected. That question distinguishes a research process from a generation process, and the distinction will matter more as the outputs of both become harder to tell apart.

WHAT THIS MEANS FOR LEADERS

- Establish what share of a file was checked against a live mail server rather than scored by a model.
- Ask whether attribution claims carry a retained, inspectable source.
- Treat pattern-inferred addresses as unverified regardless of what the process is called.
- Judge vendors on whether they report attribution and deliverability separately.

ACTION FRAMEWORK

The Verification Depth Model

Contact files sit at one of four depths, defined by what has actually been established about each record rather than by what the process is called. Each level names the specific check that unlocks the next, and the right depth for a given record depends on how that record will be used.



STAGE 01

Formatted

Addresses are syntactically valid and may have been constructed from an observed company pattern rather than observed directly. Nothing has been established about whether any mailbox exists. Files at this depth are frequently described as verified on the basis of confidence scoring. The step up is checking addresses against a live mail server rather than predicting their validity.

STAGE 02

Reachable

Domain, mailbox, and catch-all checks have been run, so deliverability is established to a high standard and hard bounces are largely eliminated. Attribution is untested, which on the modeling here leaves roughly a fifth of the file deliverable and wrong. The step up is confirming, from a source outside the mail system, that the named person holds the stated role.

STAGE 03

Attributed

Employment and role have been confirmed against a retained source, whether by a researcher or by an auditable automated process. Records are both reachable and true as at the date of checking, and the usable-record rate approaches the ceiling for a point-in-time file. The step up is not a deeper check but a repeated one, since attribution begins decaying immediately.

STAGE 04

Maintained

Verification depth is matched to intended use, recorded as a field, and enforced at routing, so records are not promoted into individual working without an upgrade check. Attributed records are re-verified on a cadence set by segment volatility, and both deliverability and attribution accuracy are measured and reported separately rather than combined into one figure.

METHODOLOGY

Accuracy figures are modeled from aggregate verification outcomes observed across Database Providers verification operations and are directional rather than measured. Deliverability means the address accepts mail at a monitored mailbox. Attribution means the named person holds the stated role at the stated organisation at the date of checking. A record is fully usable only if both are true. Modeled values are 92% deliverability and 77% attribution among deliverable records for automated stacks, giving 71% fully usable; and 97% deliverability and 97% attribution among deliverable records for human-assisted verification, giving 94% fully usable.

The composition figures in Chapter 2 decompose an automatically verified file as 71% correct and usable, 12% deliverable with a changed role, 8% undeliverable, 6% deliverable at a shared or role alias, and 3% deliverable at a generated address accepted by a catch-all domain. These sum to 100% by construction.

Cost modeling uses indicative prices of USD 0.004 per record for automated verification and USD 0.028 per record for human-assisted verification, and a handling cost of USD 3.10 per unusable record carried over from earlier reports in this series, where it is decomposed as 3.5 minutes of sales development time at a loaded rate of USD 42 per hour plus USD 0.65 of amortized platform and administration cost. Per thousand records this gives USD 903 total for automation with 710 usable contacts and USD 214 for human-assisted verification with 940 usable contacts. The break-even handling cost of approximately USD 0.07 per unusable record is derived algebraically from the same model.

The tiered case in Chapter 4 applies human verification to the 25% of records routed to individual sales development handling and automation to the remainder, with per-record handling cost applied only to the human-verified tier on the assumption that low-touch records do not consume individual attention. This gives USD 57 per thousand records for 768 usable contacts. Vendor diligence figures in Chapter 5 represent the modeled share of vendors able to answer each question concretely and are the least precise figures in the report.

NOTES AND LIMITATIONS

- All figures described as modeled are directional estimates rather than measured results from a controlled study, and will not match the experience of any individual organisation.
- Prices per record are indicative of observed market ranges and are not quotations from any named vendor.
- The cost model excludes deliverability damage from bounces and coverage effects in account-based programmes, both of which would increase the modeled cost of automated-only verification. The figures should be read as conservative.
- The tiered result depends on low-touch records genuinely not consuming per-record human attention. Where routing is not enforced, the advantage does not hold.
- Descriptions of vendor processes in Chapter 5 are general categories rather than characterisations of any named vendor.
- Database Providers uses human confirmation as part of its verification process, and readers should weigh the findings accordingly. Every cost input is stated so the comparison can be rerun independently, and Chapter 4 concludes that most of a typical file does not require human verification.

HOW DATABASE PROVIDERS CAN HELP**Ask us for the two numbers, not the one.**

The diligence questions in Chapter 5 are meant to be used on us as well as on anyone else. The one that separates methods fastest is whether a supplier can report deliverability and attribution as separate figures for the specific file you are considering, and state the date each layer was last run against it.

- Multi-layer verification combining protocol-level checks with human confirmation of employment and role, with the date of each layer stated.
- Verification of existing CRM records at the depth your downstream routing justifies, rather than one depth applied to the whole file.
- Free sample data in your target vertical and seniority band, so you can test both deliverability and attribution independently before committing.

ABOUT DATABASE PROVIDERS

Database Providers is a USA-based B2B data services company headquartered in Fremont, California. For more than a decade, the firm has helped over 750 organizations across 120+ business verticals connect with verified decision-makers through segmented email databases, data appending, verification, and enrichment services spanning healthcare, technology, manufacturing, banking and financial services, education, and real estate markets in the USA, India, the UAE, Singapore, and Malaysia.

Database Providers publishes benchmarks and practitioner guidance drawn from its verification operations, which maintain a global database of 140 million+ business contacts. Reports in this series are intended to be usable by practitioners rather than promotional, and are published with their methodology, inputs, and limitations stated in full.

10+

YEARS

750+

CUSTOMERS

120+

VERTICALS

140M+

CONTACTS

Request a free data sample

Test both numbers yourself. Request a sample in your target vertical and seniority band, measure its deliverability and its attribution separately, and compare the result against the benchmarks in this report.

DATABASE PROVIDERS

39109 Guardino Dr, Fremont, CA 94538, USA

sales@thedatabaseproviders.com

www.thedatabaseproviders.com

© 2026 Database Providers. All rights reserved. This report is provided for general informational purposes only and does not constitute legal, financial, or professional advice. Figures described as modeled or directional are estimates derived from Database Providers aggregate verification data and may not reflect the experience of any individual organization. Per-record prices are indicative of observed market ranges and are not quotations. Descriptions of verification approaches are general categories and do not characterise any named vendor. Reproduction with attribution is permitted.