Database Providers

RESEARCH REPORT No. 02

Verified vs. Scraped Data

The true cost of bad B2B contact lists, measured per usable contact rather than per record

2026 EDITION

Research by Database Providers

DELIVERABLE ACCURACY AT POINT OF PURCHASE

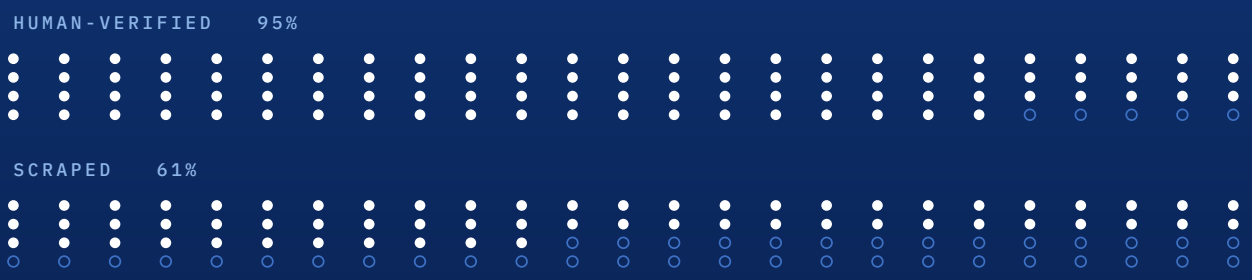

5.2x

more expensive per usable contact, despite costing 82% less per record to acquire

MODELED, AGGREGATE VERIFICATION DATA

ABOUT THIS REPORT

This report compares B2B contact lists by how they were built. It measures the starting accuracy of five acquisition methods, from unattributed scraping through multi-layer verification, and models what each one actually costs once undeliverable records, sales handling time, and deliverability damage are counted rather than assumed away.

The analysis covers deliverable accuracy at the point of purchase, hard bounce and silent-failure rates, cost per usable contact, the share of records carrying documented provenance, and the accuracy contribution of each individual verification layer. It also sets out a due-diligence checklist buyers can apply to any vendor before purchase.

Where the 2026 B2B Data Decay Report examined what time does to a contact database, this report examines what acquisition method does to it on day one. The two effects compound: a list that starts at 61% accuracy and is then held for a year without re-verification retains under half the reach its record count implies.

PUBLISHED July 2026

EDITION 2026 Edition

RESEARCH BY Database Providers

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

REPORT CODE DP-WP-2026-02

CONTENTS

FOREWORD

Every list looks the same in a spreadsheet. Two files arrive as CSVs, both with 10,000 rows, both with a name column, a title column, a company column, and an email column. One costs 400 dollars and one costs 2,200. Nothing visible in either file explains the difference, and the buyer is asked to accept on faith that the expensive one is worth five and a half times the cheap one.

This report exists because that faith is unnecessary. The difference between a scraped list and a verified one is measurable, and it is measurable in the only unit that matters to a marketing budget: cost per contact you can actually reach. Once the calculation is done that way, the pricing inverts. The cheap file becomes the expensive one, by a wide margin, and the margin is wide enough that it should change purchasing decisions rather than merely inform them.

We want to be careful about what we are and are not claiming. We sell verified data, so a report from us concluding that verified data is better invites obvious scepticism. We have tried to earn the reader's trust in two ways. First, every cost figure is built from a stated model with stated inputs, so a reader who disagrees with our assumptions can substitute their own and rerun it; the arithmetic is simple and we show it. Second, we report where verification stops helping. The gap between good verification and exhaustive verification is small, and paying for the most elaborate option available is usually not the rational choice. A report that found nothing but support for the most expensive product on our own price list would not deserve to be believed.

The finding we would most like readers to take away is not about vendors at all. It is that cost per record is the wrong unit, and that switching to cost per usable contact changes which option is cheapest. That switch costs nothing to make and it can be made before talking to any supplier, including us.

Research Team, Database Providers
Fremont, California

EXECUTIVE SUMMARY

Contact lists are priced per record and evaluated per record, but they are consumed per usable contact. This report models the gap between those two units across five acquisition methods, and finds that the cheapest list to buy is consistently the most expensive to use.

FINDING 01

Scraped lists arrive at 61% deliverable accuracy against 95% for human-verified lists.

Of a scraped file, a modeled 24% hard bounces and a further 15% resolves to addresses that accept mail without reaching a person. Human verification removes most of both categories before delivery. The difference exists at the moment of purchase, before any decay begins.

FINDING 02

Scraped data costs 5.2 times more per usable contact while costing 82% less per record.

Modeled at a 10,000-record scale, a scraped file works out to USD 2.05 per usable contact against USD 0.39 for a verified file, once undeliverable records and the sales time spent on them are counted. The acquisition discount is real and it is smaller than the handling penalty it creates.

FINDING 03

Choosing scraped over verified adds a modeled USD 8,740 in cost per 10,000 records.

The saving on purchase price is approximately USD 1,800 per 10,000 records. The additional handling cost created by undeliverable and unreachable records is approximately USD 10,540. The net position is negative before any deliverability consequences are counted.

FINDING 04

Hard bounce rates on scraped lists run near 24%, roughly thirteen times the verified rate.

This matters beyond the wasted sends, because mailbox providers assess sender reputation at domain level. A single campaign to a scraped list can suppress inbox placement for every subsequent campaign from that domain, including messages to entirely accurate records.

FINDING 05

Only a modeled 4% of scraped records carry documented provenance.

Verification confirms that an address is live. It cannot reconstruct where the record came from or on what basis it may be processed. Data protection frameworks that impose accuracy and record-keeping obligations are not satisfied by a verification pass applied after acquisition.

KEY FINDINGS AT A GLANCE

Six numbers that summarise the research findings.

61%

ACCURACY OF SCRAPED LISTS

Modeled deliverable accuracy at purchase, against 95% for human-verified lists. See Chapter 1.

5.2x

COST PER USABLE CONTACT

Modeled multiple for scraped against human-verified data at a 10,000-record scale. See Chapter 2.

82%

CHEAPER PER RECORD TO BUY

The acquisition discount on scraped data, and the reason the true cost stays hidden. See Chapter 2.

$8,740

ADDED COST PER 10,000

Modeled net penalty for choosing scraped over verified at a 10,000-record scale. See Chapter 2.

24%

HARD BOUNCE RATE, SCRAPED

Against a modeled 1.8% for human-verified data, with domain-level consequences. See Chapter 3.

4%

HAVE DOCUMENTED PROVENANCE

Share of scraped records with a recorded source and lawful basis. See Chapter 4.

CHAPTER

01

Scraped and verified lists are not the same product

Five acquisition methods, five starting accuracy levels, and one spreadsheet that hides the difference

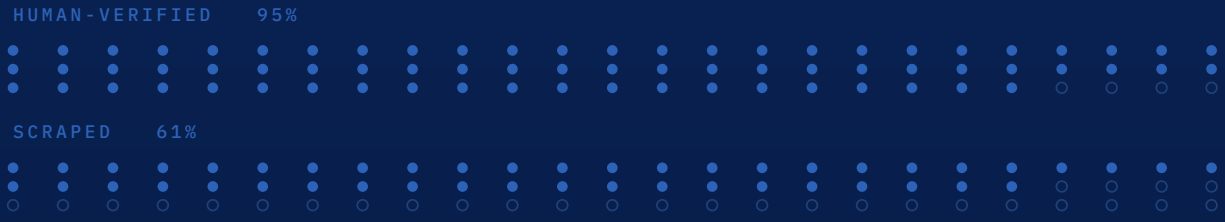

CHAPTER 01

Scraped and verified lists are not the same product

The phrase contact list describes an output format, not a product. Two files with identical columns can be built by processes with almost nothing in common, and the process determines what share of the rows correspond to a person who can be reached. That share is the only meaningful quality measure, and it is invisible in the file itself.

Five methods account for most of the B2B market. Scraping harvests addresses from public web sources, directories, and social profiles with no confirmation that they are current or monitored. Compilation aggregates published directories, registration data, and third-party files, applying format checks but not delivery checks. Permission-based collection captures addresses at the point where a person supplies them. Human verification adds a confirmation step in which a researcher establishes that the person holds the stated role at the stated organization. Multi-layer verification combines automated technical checks with that human confirmation and a defined re-verification cycle.

Ranking these by modeled deliverable accuracy at the point of purchase produces a spread of thirty-six percentage points between the extremes. The spread is wider than most buyers expect, and critically, it is present on day one. It is not a decay effect and it cannot be recovered by using the list quickly.

The gap between the weakest and strongest acquisition method is 36 percentage points, present at the moment of purchase.

Source: Database Providers, modeled deliverable accuracy at point of purchase by acquisition method, 2026.

What the missing 39% consists of

A scraped file's shortfall divides into two unequal parts, and they fail in different ways. Roughly 24% of records hard bounce, producing an immediate error the sending platform records and suppresses. These are the visible failures, and while they are damaging they are at least honest about being broken.

The remaining 15% is more troublesome. These addresses accept mail without delivering it to a person: shared aliases nobody monitors, role addresses left active after a departure, and catch-all domains that accept everything sent to them regardless of whether the local part exists. Nothing bounces. The platform reports delivery. The message reaches no one.

This second category is why delivery rate is a poor quality proxy. A scraped list can report a 76% delivery rate that looks tolerable, when the share actually reaching a human is 61%. The fifteen-point difference sits entirely inside the delivered column, depressing open and reply rates in a way that looks like a messaging problem and is not.

> *A scraped list can report a tolerable delivery rate while fifteen percent of those delivered messages reach nobody at all.*

Why scraped data starts wrong rather than becoming wrong

Scraping captures whatever a public source last published, and public sources are updated on their own schedule rather than on the world's. A directory listing may reflect a role held three years ago. A conference attendee page is accurate for the week of the conference. A profile that has not been touched since a job change presents the previous employer as current.

Pattern generation compounds this. Where a scrape finds a name and a company but no address, many tools construct one from the organization's observed format, producing an address that is plausible, correctly formatted, and frequently fictional. These constructed records are the largest single contributor to catch-all acceptance, because a catch-all domain accepts a fabricated local part exactly as readily as a real one.

Neither problem is a decay problem. Both are present the moment the file is delivered, which is why the remedy is verification at source rather than a faster campaign schedule.

WHAT THIS MEANS FOR LEADERS

- Ask any vendor for deliverable accuracy at point of purchase, not total record count.
- Treat delivery rate as an incomplete quality measure and ask for the accept-all share separately.
- Establish whether any addresses in a file were pattern-generated rather than observed.
- Do not expect fast campaign execution to compensate for low starting accuracy.

CHAPTER

02

The price you pay is the smallest cost you incur

Switching from cost per record to cost per usable contact inverts the pricing of every option

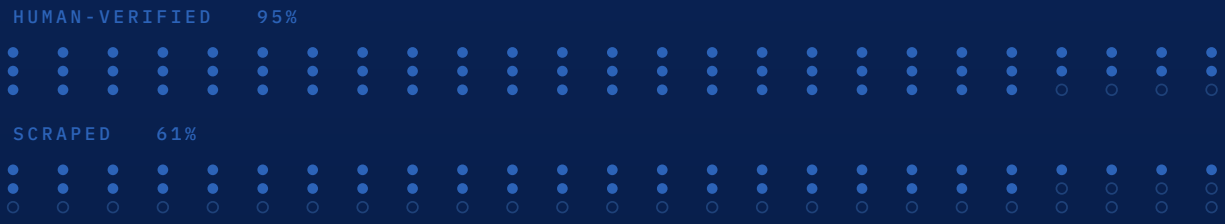

CHAPTER 02

The price you pay is the smallest cost you incur

List pricing is quoted per record because record count is easy to verify and accuracy is not. The unit favours the seller of low-accuracy data, since it prices the rows rather than the reachable contacts inside them. Restating the same offers as cost per usable contact requires two inputs a buyer already has: the price, and an honest estimate of the deliverable share.

The model applied here is deliberately simple so that readers can substitute their own numbers. It assumes a 10,000-record purchase and charges each unusable record with the sales development handling it consumes before it is identified as dead. That handling is modeled at 3.5 minutes at a loaded rate of USD 42 per hour, or USD 2.45, plus USD 0.65 of amortized platform and administration cost, giving USD 3.10 per unusable record.

Applied across the five acquisition methods, the ranking reverses completely. The cheapest file to buy becomes the most expensive to use, and the difference is not marginal.

Restated per usable contact, scraped data is the most expensive option on the market by a factor of five.

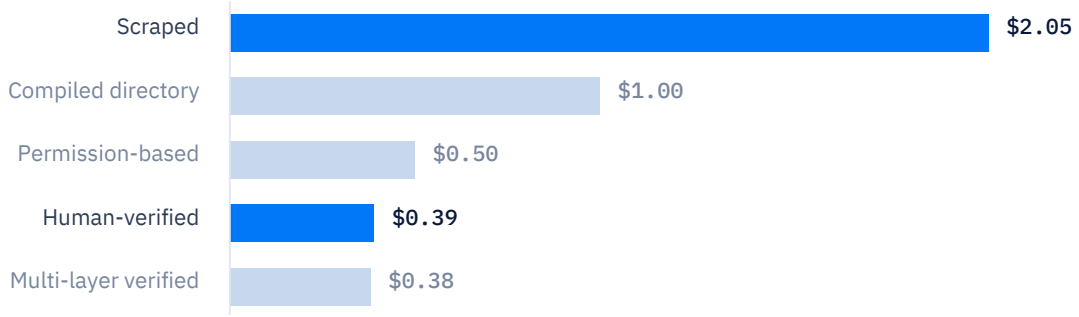

Source: Database Providers, modeled cost per usable contact at 10,000-record scale, 2026.

The calculation in full

For a scraped file at USD 0.04 per record, the purchase costs USD 400. At 61% accuracy, 6,100 records are usable and 3,900 are not. The unusable records consume 3,900 multiplied by USD 3.10, or USD 12,090, in handling. Total cost is USD 12,490 for 6,100 usable contacts, which is USD 2.05 each.

For a human-verified file at USD 0.22 per record, the purchase costs USD 2,200. At 95% accuracy, 9,500 records are usable and 500 are not, consuming USD 1,550 in handling. Total cost is USD 3,750 for 9,500 usable contacts, or USD 0.39 each.

The scraped file saves USD 1,800 at purchase and adds USD 10,540 in handling, a net penalty of USD 8,740 per 10,000 records. At 100,000 records the modeled penalty is USD 87,400. Both figures exclude deliverability damage, which Chapter 3 addresses and which no per-record model captures.

A buyer who disputes the USD 3.10 handling figure can solve for the break-even point directly. Scraped and verified data reach parity when an unusable record costs approximately USD 0.28 to handle, which is about 24 seconds of sales development time at the modeled rate. Above that, verified data wins on cost alone. The figure is worth pausing on, because USD 0.28 is lower than the amortized platform and administration cost of simply carrying a record in a CRM. Verified data therefore comes out ahead under essentially any assumption a buyer might substitute, including one where sales development time is treated as free.

> *Scraped and verified data break even at 24 seconds of attention per dead contact. No dead contact costs 24 seconds.*

Where paying more stops helping

The same model that condemns scraped data also sets a ceiling on how much verification is worth buying. Multi-layer verification at 97% accuracy costs a modeled USD 0.38 per usable contact against USD 0.39 for human verification at 95%. The difference is a single cent, and it sits well inside the error bars of the model.

The practical reading is that the large gains come from crossing the line into verification at all, not from buying the most exhaustive tier available. Organizations sending at very high volume, or into markets where bounce tolerance is unusually low, have a legitimate case for the top tier because the deliverability protection compounds with volume. For most buyers, standard human verification captures nearly all of the available benefit.

We state this plainly because the opposite claim would be commercially convenient for us and would not survive the arithmetic. The honest conclusion is that verification is worth buying and premium verification is worth buying only under specific conditions.

WHAT THIS MEANS FOR LEADERS

- Re-quote every list offer as cost per usable contact before comparing prices.
- Ask vendors to state deliverable accuracy contractually so the calculation can be made.
- Use your own loaded hourly rate in the model rather than an industry average.
- Do not pay for the most exhaustive verification tier without a volume or deliverability reason.

CHAPTER

03

Bounce damage is domain-level and outlives the campaign

The cost that no per-record model captures, and the reason a single bad send has a long tail

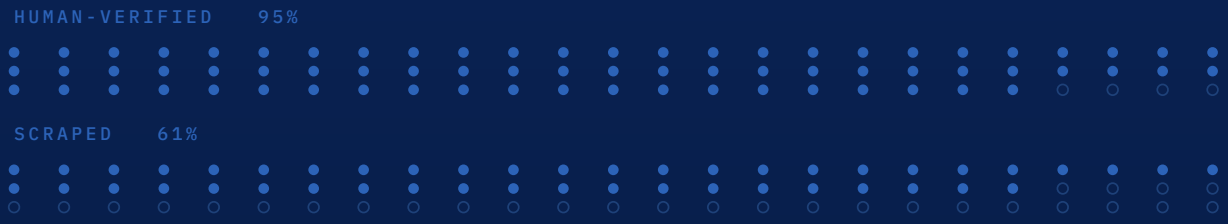

CHAPTER 03

Bounce damage is domain-level and outlives the campaign

Cost per usable contact understates the penalty for bad data, because it treats each undeliverable record as a self-contained loss. Mailbox providers do not. Sender reputation is assessed at the level of the sending domain and the sending infrastructure, which means the consequences of one campaign attach to every campaign that follows from the same domain.

Modeled hard bounce rates across acquisition methods span an order of magnitude. Scraped files bounce at roughly 24%, compiled files at 13%, permission-based at 5.5%, human-verified at 1.8%, and multi-layer verified at 1.1%. General industry guidance treats sustained hard bounce rates above about 2% as a reputation risk, which places three of the five methods outside tolerance on their own.

Three of the five acquisition methods bounce above the level generally treated as a reputation risk.

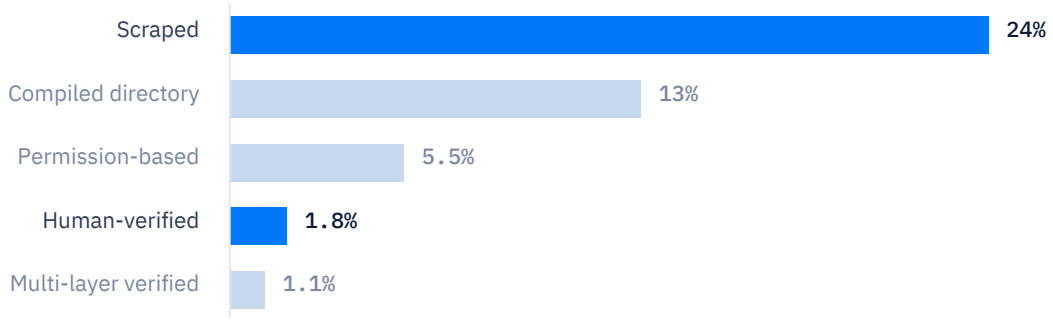

Source: Database Providers, modeled hard bounce rate by acquisition method, 2026.

The authentication baseline is now a floor, not an advantage

Sender requirements have tightened materially. Google and Yahoo introduced bulk sender requirements effective February 2024. Microsoft announced comparable requirements for its consumer domains in April 2025, with enforcement beginning on 5 May 2025 for domains sending more than 5,000 messages per day to Outlook.com, Hotmail.com, and Live.com. The shared baseline across providers includes a valid SPF record, DKIM signing, a published DMARC policy of at least p=none aligned with SPF or DKIM, functional and reply-capable sender addresses, and a working unsubscribe mechanism. Google additionally publishes a spam complaint rate threshold of 0.3% for bulk senders.

Authentication and list quality are independent prerequisites, and neither substitutes for the other. Correct authentication does not protect a sender using a scraped list, because bounce and complaint signals are generated by the recipients rather than the DNS configuration. A clean list does not protect a sender with broken authentication, because messages may be rejected before content is assessed.

The change worth noting is that authentication used to be a differentiator and is now a floor. Because most legitimate senders now comply, list quality has become the variable that separates senders who reach the inbox from senders who do not.

Spam traps and the cost of a single hit

Scraped lists carry a specific hazard that compiled and verified lists largely avoid. Spam traps are addresses published deliberately in places only a harvester would find them, or recycled from long-abandoned mailboxes, and they exist to identify senders who acquire addresses without consent. Because scraping is precisely the behaviour traps are designed to detect, scraped files contain them at materially higher rates than any other acquisition method.

A trap hit is not scored like a bounce. Depending on the operator, it can result in immediate listing of the sending IP or domain on a blocklist, which affects delivery to every provider consulting that list rather than only the one whose trap was hit. Recovery is a process rather than a setting change: delisting requests, demonstrated remediation, and a period of clean sending before reputation recovers, commonly running to about 90 days.

Costed against the model in Chapter 2, a single blocklisting that suppresses outbound performance for a quarter dwarfs the entire acquisition saving on a scraped list. The asymmetry is the point. The saving is bounded at roughly USD 1,800 per 10,000 records; the downside is not bounded at all.

> *The saving on a scraped list is bounded at eighteen hundred dollars. The downside is not bounded.*

WHAT THIS MEANS FOR LEADERS

- Treat any list projected to bounce above 2% as a domain-level risk, not a campaign-level one.
- Verify SPF, DKIM, and DMARC alignment independently of list quality, since both are prerequisites.
- Use a separate subdomain for cold outbound so reputation damage cannot reach transactional mail.
- Price the tail risk of a blocklisting into any decision to buy unverified data.

CHAPTER

04

Scraped data carries provenance risk no verification pass can retrofit

Verification establishes that an address works, not that it may lawfully be used

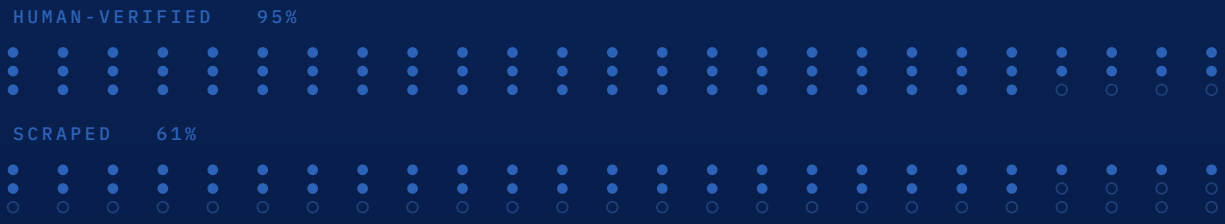

CHAPTER 04

Scraped data carries provenance risk no verification pass can retrofit

There is a category difference between two questions a buyer needs answered. The first is whether an address reaches the intended person, which verification answers well. The second is where the record came from and on what basis it may be processed, which verification does not answer at all. A verification pass applied to a scraped file produces a file of confirmed-live addresses with unknown origins.

Documented provenance means a recorded source for each record, the date it was obtained, and the basis on which it is held. Modeled across acquisition methods, the share of records carrying that documentation ranges from a modeled 4% for scraped data to 96% for permission-based collection, where the collection event itself creates the record.

Permission-based collection documents provenance by construction; scraping documents almost nothing.

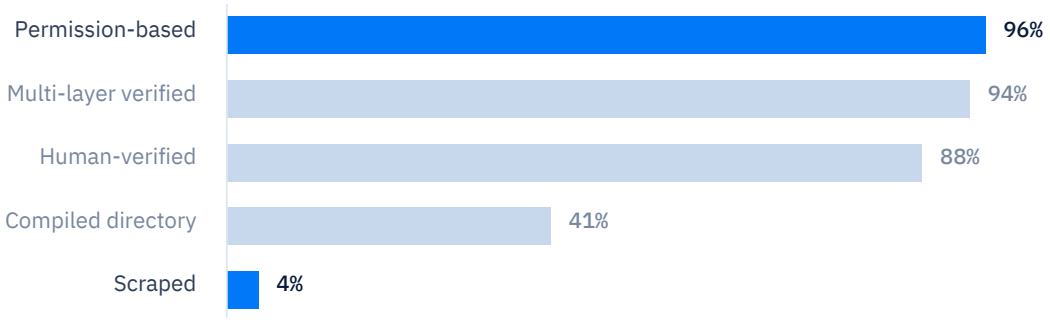

Source: Database Providers, modeled share of records with documented source and basis, 2026.

Why the documentation is the asset

Several data protection frameworks impose obligations that are discharged through records rather than through outcomes. Accuracy obligations require that personal data be kept accurate and up to date. Transparency obligations require that a data subject can be told where their data came from. Erasure and objection rights require that a request can be honoured and propagated across every copy held. Each of these is straightforward with documented provenance and close to impossible without it.

The frameworks in scope for a business selling across the markets Database Providers covers include the European Union General Data Protection Regulation, in force since 2018 and the structural reference for much of what followed; India's Digital Personal Data Protection Act; the federal personal data protection law in the United Arab Emirates; Singapore's Personal Data Protection Act; Malaysia's Personal Data Protection Act; and in the United States a sectoral and state-level patchwork including the California Consumer Privacy Act as amended.

The practical exposure is not usually a regulator. It is a single subject access or erasure request that cannot be answered. A buyer holding a scraped file who receives a request asking where the data came from has no answer available, and no verification vendor can supply one after the fact, because the information was never captured.

Due diligence a buyer can actually run

Provenance is checkable before purchase, and the questions are short enough to put in an email. Ask for the source category of the records, the date range over which they were obtained, the lawful basis relied on for the

markets in scope, and whether the vendor will state deliverable accuracy contractually. Ask what happens to a record when an erasure request is received, and whether suppression persists across future refreshes of the same file. Ask for a sample and test it independently rather than accepting a quality claim.

The pattern in the answers matters more than any single response. Vendors who maintain provenance records answer these questions concretely and quickly, because the information is already recorded. Vendors who do not tend to answer with reassurance rather than specifics, redirect to compliance certifications that do not address origin, or offer to run a verification pass as though it resolved the question.

None of this constitutes legal advice, and obligations differ materially by market and by the basis on which data is held. Organizations operating across several of these jurisdictions should confirm their position with qualified counsel rather than relying on a vendor's characterisation, including ours.

WHAT THIS MEANS FOR LEADERS

- Require documented source, acquisition date, and lawful basis as a condition of purchase.
- Confirm that suppression and erasure records persist across every future file refresh.
- Judge vendors on whether they answer provenance questions specifically or reassuringly.
- Do not treat a verification pass as a remedy for unknown data origin.

CHAPTER

05

A verification stack is auditable; a scrape is not

What each verification layer contributes, and how much of the gain arrives before the expensive step

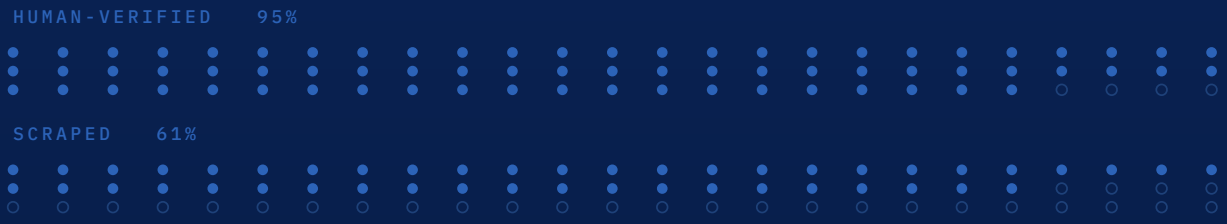

CHAPTER 05

A verification stack is auditable; a scrape is not

Verification is often sold as a single undifferentiated service, which makes it hard to evaluate and easy to overpay for. It is more usefully understood as a sequence of independent layers, each removing a different failure mode and each contributing a measurable share of the total accuracy improvement.

Syntax validation checks that an address is correctly formed. Domain and MX validation confirms the domain exists and is configured to receive mail, removing records whose employers have retired a domain. SMTP validation opens a conversation with the receiving server to establish whether the specific mailbox exists. Catch-all detection identifies domains that accept all mail regardless of local part, which is what separates a real delivery from an apparent one. Human confirmation establishes that the person holds the stated role at the stated organization, which is the only layer that validates the record's meaning rather than its mechanics.

Measured cumulatively from a raw scraped base, the layers produce a rising curve with the largest single jump at SMTP validation and the final correction at human confirmation.

Automated layers carry accuracy from 68% to 92%; only human confirmation closes the last five points.

Source: Database Providers, modeled cumulative accuracy by verification layer applied to a raw base, 2026.

Why the last five points cost the most

The automated layers are cheap because they are mechanical. They run at scale, cost fractions of a cent per record, and require no judgement. They also share a blind spot: every one of them tests whether an address works, and none tests whether the person behind it is who the record claims.

That blind spot is where the residual eight percent lives after catch-all detection. A live, correctly configured, monitored mailbox belonging to someone who left the company last year passes every automated check. So does a live mailbox belonging to someone who was promoted out of the role that made them a relevant prospect. The record is technically valid and commercially useless, and only a human check catches it.

This is also why verification cadence matters more than verification depth, a point the 2026 B2B Data Decay Report develops at length. Role accuracy is the component that decays, and re-verifying quarterly with a standard stack outperforms verifying once with an exhaustive one.

The auditability difference

The structural argument for a verification stack is not only that it produces better data. It is that each layer generates a record of what was checked and when, which makes the resulting file auditable. A buyer can ask which layers were applied, on what date, and to what share of the file, and receive an answer that can be tested

against a sample.

A scrape generates no such trail by construction. There is no checkpoint to interrogate, no date attached to a confirmation that never happened, and no way to distinguish a record that was observed from one that was pattern-generated. This is what makes the two products different in kind rather than in degree: one can be evaluated and the other can only be trusted.

For buyers, the practical translation is a single question that separates the categories quickly. Ask a vendor to state, for a specific sample, which verification layers were applied and on what date each was last run. The answer, or the absence of one, is more informative than any accuracy claim on a marketing page.

> *One product can be evaluated and the other can only be trusted. That is a difference in kind, not degree.*

WHAT THIS MEANS FOR LEADERS

- Ask which specific verification layers were applied and when each was last run.
- Expect catch-all detection explicitly, since it is what distinguishes real delivery from apparent delivery.
- Prioritise verification cadence over verification depth, because role accuracy is what decays.
- Test any accuracy claim against an independently checked sample before committing to volume.

ACTION FRAMEWORK

The Contact Data Assurance Model

Contact data can be positioned on four levels of assurance, defined by what a holder is able to demonstrate about it rather than by what it cost. Each level is a claim a buyer can either evidence or cannot, which makes the model usable as a procurement standard rather than only as a description.

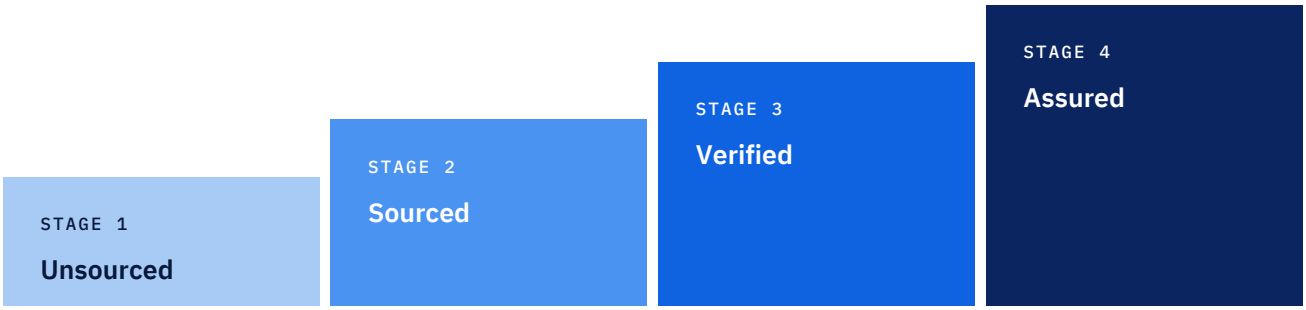

STAGE 01

Unsourced

The holder cannot say where records came from, when they were obtained, or whether any address was observed rather than generated. Quality is asserted rather than evidenced, and no erasure or transparency request can be answered fully. Moving up does not require better data; it requires beginning to record source and date at the point of acquisition, which is available to any buyer immediately.

STAGE 02

Sourced

Source category and acquisition date are recorded for every record, and suppression lists persist across refreshes. Accuracy is still unmeasured, so cost per usable contact cannot be calculated and campaign performance remains unpredictable. The step up is technical validation, which converts an unknown accuracy level into a measured one and makes procurement comparisons possible for the first time.

STAGE 03

Verified

Records pass technical validation including domain, mailbox, and catch-all checks, with role accuracy confirmed by human review. Deliverable accuracy is stated as a number and can be tested against a sample, so lists can be compared on cost per usable contact rather than price per record. Most of the available economic benefit is captured here. Advancing further means adding a defined cadence and an audit trail.

STAGE 04

Assured

Every record carries documented source, lawful basis, and a dated log of which verification layers were applied. Re-verification runs on a defined cadence set by segment volatility rather than by calendar convenience. Accuracy is contractual rather than promotional, erasure requests propagate automatically across all copies, and the file can be audited by a third party without reconstruction.

METHODOLOGY

Accuracy figures in this report are modeled from aggregate verification outcomes observed across Database Providers verification operations. A record is treated as usable when it reaches the intended person at the intended organization. Records are treated as unusable if they hard bounce, if they resolve to a shared or unmonitored alias, if they are accepted by a catch-all domain without a corresponding mailbox, or if the named person no longer holds the stated role at the stated organization. Acquisition method categories reflect how records entered the file rather than any vendor's description of its own process.

Cost modeling uses a 10,000-record purchase and the following stated inputs: indicative per-record prices of USD 0.04 for scraped, USD 0.10 for compiled, USD 0.18 for permission-based, USD 0.22 for human-verified, and USD 0.28 for multi-layer verified data; and a handling cost of USD 3.10 per unusable record, comprising 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. Prices are indicative of observed market ranges rather than quotations, and readers are encouraged to substitute their own figures, which changes absolute values but not the ordering of the results.

The break-even handling cost of approximately USD 0.28 per unusable record is derived algebraically from the same model by solving for the point at which scraped and human-verified total cost per usable contact are equal, and is expressed in time at the same USD 42 per hour loaded rate. Provenance figures represent the modeled share of records for which a source category, acquisition date, and stated basis are recorded. Verification layer figures are cumulative and are modeled against a raw scraped base rather than measured independently at each layer.

External requirements are drawn from published sender documentation rather than modeled. Google and Yahoo introduced bulk sender requirements effective February 2024. Microsoft announced comparable requirements for its consumer domains in April 2025, with enforcement beginning 5 May 2025 for senders exceeding 5,000 messages per day. Data protection frameworks are referenced by name and general obligation only, without interpretation of their application to any specific processing activity.

NOTES AND LIMITATIONS

— All accuracy, cost, and provenance figures are modeled directional estimates rather than measured results from a controlled study, and will not match the experience of any individual organization.
— Database Providers sells verified contact data, and readers should weigh the findings accordingly. Every cost input is stated so the model can be independently rerun with different assumptions.
— Per-record prices are indicative of observed market ranges and are not quotations from any named vendor.
— Modeled rates reflect the composition of the analyzed database and may not generalize to segments, verticals, or markets outside it.
— This report describes data protection frameworks in general terms and does not constitute legal advice. Obligations vary by jurisdiction and by the basis on which data is held; confirm your position with qualified counsel.
— Where external sender requirements are cited, senders should verify current provider documentation directly, as authentication and filtering requirements continue to change.

HOW DATABASE PROVIDERS CAN HELP

Run the cost-per-usable-contact test on your own file.

The model in Chapter 2 needs only two inputs: your list price and its true deliverable accuracy. The second is the one most buyers cannot obtain. Database Providers will verify a sample of your existing file so you can calculate it, and provide a matched verified sample for comparison.

— Verified, segmented contact databases across 120+ verticals in the USA, India, the UAE, Singapore, and Malaysia.

— Multi-layer verification of existing CRM records, including catch-all detection and human role confirmation.

— Free sample data in your target vertical, seniority band, and market, for independent accuracy testing before any commitment.

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+	750+	120+	140M+
YEARS	CUSTOMERS	VERTICALS	CONTACTS

Request a free data sample

Test the findings in this report against your own file. Request a verified sample in your target vertical, seniority band, and market, and measure its accuracy independently before you commit to anything.

DATABASE PROVIDERS

39109 Guardino Dr, Fremont, CA 94538, USA

sales@thedatabaseproviders.com

www.thedatabaseproviders.com