

THE BAKERSFIELD CYBERSECURITY REPORT · 2026 EDITION

Secure by Default, Exposed by Omission

Email authentication and web security posture among 4,406 independent businesses in Bakersfield, California

Scott Wattenbarger

Mitivanta, Bakersfield, California · scott@mitivanta.com

Data collected 2026-10-18. Population: 4,406 live, independent businesses. Source: Overture Maps, release 2026-09-23.1.

Read online, with interactive figures and the data behind each one: mitivanta.com/bakersfield-cybersecurity-report

86%

of surveyed Bakersfield
business domains lack an
enforced DMARC policy
3,790 of 4,402

1 in 17

publish a DMARC
reject policy
5.8% of domains

8.5%

pass five basic, free
security checks
357 of 4,183 fully assessed

88%

of medical, dental and
veterinary practices lack an
enforced DMARC policy
524 practices

1,117

local business listings point
to a domain with no DNS
delegation
995 with no
registration record

CITE THIS REPORT

Wattenbarger, S. (2026). The Bakersfield Cybersecurity Report 2026: Secure by Default, Exposed by Omission. Mitivanta Research. <https://mitivanta.com/bakersfield-cybersecurity-report>

Abstract

Background. Business email compromise cost US victims a reported \$3.05 billion in 2025, the second costliest category of cybercrime reported to the FBI, and much of it relies on email that appears to come from a trusted business. Three published standards, SPF, DKIM and DMARC, let a domain owner tell receiving mail servers how to treat mail that falsely uses the owner's exact domain. They carry no licence fee, though deploying them correctly takes work, and they do not address compromised accounts, lookalike domains or display name impersonation. Measurement research has documented their adoption among large and popular domains. Far less is known about the independent small businesses that make up most local economies.

Methods. We built a population of 4,406 live, independent businesses in Bakersfield, California from the Overture Maps open places dataset, excluding chains, franchises, public bodies and domains that no longer exist. For each we read public DNS and registration records, made one request to the home page, and generated typosquatting candidates (269,689 in total). We report proportions with Wilson 95% confidence intervals, compare groups with Pearson's chi-square test and Cramér's V, and fit multivariable logistic regression models. An independent implementation, run on a stratified sample of 491 businesses, agreed with 99.8% of DMARC and 100.0% of SPF classifications after two errors it exposed were corrected.

Results. 86.1% (95% CI 85.0% to 87.1%) of business domains publish no enforced DMARC policy, leaving receivers without the owner's instruction to quarantine or reject mail that fails authentication; 5.8% publish a reject policy. Of the 4,183 businesses where all five could be assessed, only 8.5% (95% CI 7.7% to 9.4%) pass five basic, free controls. Web transport security is far stronger: 93.4% negotiate TLS 1.3. Website controls largely follow platform defaults (hosted site builders send at least one security header 87.7% of the time, against 40.1% for self managed sites), while email controls depend on deliberate configuration: after adjustment, businesses on Microsoft 365 (OR 0.27) or behind a filtering gateway (OR 0.28) have roughly a quarter of the odds of lacking an enforced policy that Google Workspace customers have. Email controls cluster strongly with one another (DMARC enforcement with a published DKIM key, OR 12.0) and web controls with one another (HTTPS redirect with HSTS, OR 6.3), while the raw association across the two families is near zero (DMARC enforcement with HSTS, OR 0.91, not significant). Adjusting for website platform reveals a modest positive association (OR 1.43), masked in the raw data by hosted builders that send web headers by default but rarely enforce DMARC. Separately, 1,117 of 5,722 business map listings (19.5%) point to a domain with no DNS delegation; the registry had no record of 995 of them.

Conclusions. In this population, measured configuration appears to be shaped less by owners' decisions than by platform defaults and, we hypothesise, by who configured each system. Web transport has acquired secure defaults; email authentication has not. These are observations of public configuration, not of attacks or delivery outcomes. We discuss implications for owners, IT providers, platforms and policy.

Keywords: email authentication; DMARC; SPF; DKIM; business email compromise; small business; HTTP security headers; typosquatting; internet measurement; Bakersfield

KEY FINDINGS AT A GLANCE

- **86% of surveyed Bakersfield business domains lack an enforced DMARC policy.** Only 1 in 17 publishes a reject policy.
- **Few pass five free, basic controls** (8.5% of those fully assessed). Most fall short on email, not on their website.
- **Website controls appear to follow platform defaults; email authentication has to be configured.** Site builders apply transport protections automatically and do nothing for the email domain.
- **Email and website protections are only weakly linked.** Email controls cluster tightly with each other, and so do website controls. Across the two the raw link is near zero; a modest positive link appears once website platform is taken into account.
- **Provider type predicts enforcement.** Microsoft 365 and filtering gateway customers have about a quarter of the odds of lacking enforcement. We hypothesise this reflects who set them up; the study did not measure that directly.
- **1 in 5 map listings points to a domain with no DNS delegation.** The registry returned no registration record for 995 of the 1,117; whoever registers such a name can receive mail sent to the old addresses.

1. Introduction

One way criminals take money from small businesses is not by breaking into their systems but by pretending to be them, or someone they trust. An email arrives that appears to come from a supplier, a client, or the business's own owner, asking for an invoice to be paid to new bank details or for a quick

purchase of gift cards. In 2025 the FBI's Internet Crime Complaint Center received 24,768 complaints of business email compromise with reported losses of \$3.05 billion, the second costliest category after investment fraud, within total reported losses of \$20.9 billion [14]. The FBI separately puts exposed losses to this one scheme at more than \$55 billion between October 2013 and December 2023 [13]. The 2026 Verizon Data Breach Investigations Report finds small organisations disproportionately affected by ransomware [51].

Some of this fraud exploits a simple weakness: the global email system does not, by default, check that a message claiming to come from a domain was actually sent by that domain's owner. Three standards address this for the domain itself, though not for compromised accounts or lookalike domains. The Sender Policy Framework (SPF) lists the servers allowed to send a domain's mail [24]. DomainKeys Identified Mail (DKIM) signs messages cryptographically [8]. Domain based Message Authentication, Reporting and Conformance (DMARC) ties the two to the address a person actually sees, and tells receiving servers what to do when a message fails: deliver it, quarantine it, or reject it [19]. None carries a licence fee; all three are published as DNS records and supported by every major mailbox provider, and deploying them correctly takes care, because a mistake can block legitimate mail.

Internet measurement research has studied the adoption of these standards extensively, but almost always among the most popular domains on the internet, or among large organisations [3, 10, 12, 25, 30, 52]. Those populations are dominated by companies with security teams. The independent small business, a dental practice, a contractor, a family restaurant, a two person law office, is the typical employer in most American cities and the typical victim of invoice fraud, yet its security posture is rarely measured directly.

This paper measures it for one city. Its contributions are:

1. **A population level measurement** of email authentication, web transport security, software exposure, domain lifecycle and typosquatting exposure for 4,406 independent businesses in Bakersfield, California, built from an open dataset with a documented, reproducible inclusion process.
2. **Evidence that website controls in this population largely follow platform defaults**, while email authentication depends on deliberate configuration, and that platform defaults mask a modest positive association between the two.
3. **A multivariable model** showing that the strongest predictor of email protection is the kind of provider a business uses, which we hypothesise, without having measured it, reflects whether a professional set it up; sector and neighbourhood matter less.
4. **A measurement of residual trust risk**: the share of local business listings that point to domains with no DNS delegation, and the registry status of those domains.
5. **Practical guidance** ordered by value, for owners and for the IT providers who serve them.

2. Background and prior work

2.1 Email authentication

Email was designed without sender verification. The address in the From line that a person reads is a field the sender fills in. SPF [24] lets a domain publish, in a DNS TXT record, the list of servers permitted to send its mail. But SPF checks the envelope sender, an address hidden from the reader, so on its own it does not stop a forgery of the visible From address. SPF also has a hard limit: evaluating a record may require no more than ten DNS lookups, and a record that exceeds this produces a permanent error and protects nothing.

DKIM [8] lets a sending server sign each message with a private key whose public half is published in DNS under a named selector. A receiver can verify that the message was signed by the domain and not altered. Neither SPF nor DKIM tells the receiver what to do when a check fails.

DMARC [19, 28] supplies that instruction and connects the checks to the visible From domain through alignment. A DMARC record at `_dmarc.example.com` carries a policy: `p=none` (monitor only: the owner asks receivers not to change how they handle failing mail), `p=quarantine` (treat failures as suspicious, usually the spam folder) or `p=reject` (refuse failures). In every case the receiver makes the final decision; a DMARC policy states the domain owner's preference, and receivers also apply their own filtering [19]. Nor does any policy address every impersonation: a reject policy does nothing about mail from a compromised account, from a lookalike domain, or that fakes only the display name. It can also name an address to receive aggregate reports (`rua`), which show the owner every source sending mail in the domain's name. The intended deployment path is to begin at `p=none` with reporting, correct every legitimate sender, then move to quarantine and finally reject. DMARC was first published in 2015 as an Informational specification [28]. A revision placing it on the IETF Standards Track was published in May 2026 as RFC 9989, with aggregate reporting specified separately in RFC 9990 [5, 19]. Whether reports go anywhere is itself a signal of whether anyone manages a domain [3].

Three outside pressures have pushed organisations toward DMARC since 2023. **Mailbox providers:** since February 2024, Google and Yahoo have required senders of more than 5,000 messages a day to their users to authenticate with SPF and DKIM and to publish a DMARC record, although a monitoring policy of `p=none` satisfies the requirement [16, 29, 56]; Microsoft announced equivalent requirements for Outlook.com in April 2025, enforced from May 2025 [45]. **Card payments:** PCI DSS v4.0.1 requirement 5.4.1, mandatory since 31 March 2025, calls for processes and automated mechanisms to detect and protect personnel against phishing, and its guidance names DMARC, SPF and DKIM as examples of anti-spoofing controls [43]. **Government guidance:** joint guidance from CISA, the NSA, the FBI and MS-ISAC, which includes a section for small and medium businesses, advises setting DMARC to reject [9], alongside NIST's Trustworthy Email guidance [47] and the Cybersecurity Framework 2.0 with its small business quick start guide [39, 40].

Two further standards protect mail in transit rather than its origin. MTA-STS [36] lets a domain require encrypted delivery, and SMTP TLS reporting [35] reports failures. They are rarely deployed outside large organisations and are reported here for completeness.

2.2 Web transport and browser protections

HTTPS protects traffic between a visitor and a website, using TLS, now in version 1.3 [46]. The free, automated certificate authority Let's Encrypt [1] transformed HTTPS adoption in the last decade. A site can still fail its visitors in two common ways: by serving its plain http address without redirecting to https, and by presenting a certificate that is expired or does not cover the address typed. HTTP Strict Transport Security (HSTS) [20] instructs browsers to use only the encrypted site from then on, which prevents a downgrade on later visits [26]. Further response headers restrict what a page may do: Content Security Policy limits where scripts may load from [53], frame protection prevents the page being embedded to trick clicks, and others govern content sniffing, referrer leakage and access to device features. Adoption across the web remains partial: the HTTP Archive's 2025 Web Almanac reports a Content Security Policy on 21.9% of pages and HSTS on 36% of mobile pages [50]. Certificate management is also about to become harder to do by hand. In April 2025 the CA/Browser Forum adopted Ballot SC081v3, which reduces the maximum validity of publicly trusted TLS certificates from 398 days to 200 days for certificates issued from 15 March 2026, 100 days from 15 March 2027, and 47 days from 15 March 2029 [6].

2.3 Software exposure

Many small business websites run a content management system, most often WordPress, that the owner is responsible for updating. Sites frequently announce their software and version in page markup or response headers. Outdated client side libraries are common on the web [31]: a longitudinal study found that an average of 41.2% of websites carry at least one vulnerable client side resource, with an average window of vulnerability of 531 days [33]. jQuery releases before 3.5.0 are affected by two cross site scripting flaws, CVE-2020-11022 and CVE-2020-11023 [38]. A visible version number does not prove a site is exploitable, but it is exactly what automated scanning tools read first. In the wider WordPress ecosystem, Patchstack recorded 11,334 new vulnerabilities in 2025, 91% of them in plugins rather than in WordPress itself [42].

2.4 Domain lifecycle and lookalike domains

A domain is leased, not owned. When a registration lapses and the registry releases the name, it can be registered by anyone, who then inherits the trust others still place in it: email sent to old addresses, links in old listings and documents [32, 48]. In one experiment, 201 re-registered domains received more than 650 million requests in four months [48]. Separately, criminals register typosquatting domains, names one plausible typo away from a real one [2, 49], and an analysis of 690,502 phishing domains found 66.1% had been registered for the purpose [34]. A lookalike domain with mail records configured can send email that passes every authentication check, because the criminal really does own it. The Registration Data Access Protocol (RDAP) [21] provides structured public registration data, and CAA records [17] and DNSSEC let owners restrict certificate issuance and sign their DNS.

2.5 Related work

Durumeric and colleagues measured email delivery security across the internet and found widespread gaps in both transport encryption and sender authentication [12]. Foster and colleagues showed that the protection a user enjoys depends heavily on the provider [15], a theme this paper finds again from the sender's side. Hu and Wang demonstrated end to end that spoofed email reaches inboxes when sender

domains lack enforced policies [22]. More recent work measures each standard's configuration at scale. Czybik and colleagues found SPF on 56.5% of 12 million domains, with errors in 2.9% of records and over a third permitting mail from more than 100,000 addresses [10]. Wang and colleagues measured DKIM deployment across the most popular million domains [52], Ashiq and colleagues examined DMARC reporting and its security implications [3], and Hureau and colleagues analysed what holds organisations back from deploying DMARC [23]. Two studies examine the period after the 2024 sender requirements: Kondo and colleagues measured which domains had not implemented the measures [25], and Länge and colleagues, scanning the 100,000 most popular domains daily for eight months from April 2025, found DMARC adoption rising from 55.6% to 58.4%, while only 13.1% to 14.8% of domains combined a reject policy with reporting; a monitoring policy with reporting was the most common configuration [30]. Krombholz and colleagues showed that even experienced administrators find correct HTTPS deployment difficult [27], and Herley argued that users rationally ignore security advice whose cost exceeds its visible benefit [18]. Studies of typosquatting have documented its scale and persistence [2, 49], and Lever and colleagues showed that lapsed domains retain exploitable residual trust [32]. To our knowledge, no published study has measured these properties across an enumerated local population of independent small businesses.

2.6 What this study adds

The measurement techniques here are not new: reading SPF, DMARC and DKIM records from DNS, registration data from RDAP, and response headers from a single request are the established methods of the work above, and the broad finding that email authentication is unevenly deployed is well documented. Table 1 sets the closest prior work beside this study. What this study adds is in its population and its design:

1. **The population.** Prior measurements draw on popularity ranked lists of domains, on all registered domains, or on crawls of the web. Those populations are dominated by large organisations, or mix businesses with personal, parked and abandoned domains. This study enumerates one city's independent businesses from an open places dataset, with a documented inclusion process (Figure 1), so that its findings describe the small local business specifically.
2. **Email and web measured on the same businesses.** Prior studies measure email authentication or web security, each on its own population. Measuring both families of controls on the same 4,406 businesses is what allows the association between them to be estimated (Section 5.9), including the finding that platform defaults mask a modest positive association.
3. **Organisational covariates.** Sector, neighbourhood, email provider, website platform and domain age are recorded for every business, so variation in configuration can be modelled rather than only described (Section 5.7).
4. **Listings as an exposure.** Residual trust in lapsed domains has been measured by re-registering them [48]; this study measures how many of the domains named in one city's business listings no longer have a DNS delegation, and their registry status (Section 5.8).

The comparison cuts both ways. Prior studies cover far larger populations and, in several cases, longer periods or end to end tests of delivery that this study does not attempt; this study covers one city on one day. Its contribution is local evidence about a population that large scale studies do not isolate, not a new method or a general estimate for small businesses everywhere.

Table 1. This study beside the closest prior work

STUDY	POPULATION	WHAT WAS MEASURED
Durumeric and colleagues, 2015 [12]	Internet scale: popular domains and mail servers	Email delivery security: transport encryption and sender authentication
Hu and Wang, 2018 [22]	Popular email providers	End to end delivery of spoofed email
Wang and colleagues, 2022 [52]	The most popular million domains, over time	DKIM deployment and misconfiguration
Czybik and colleagues, 2023 [10]	12 million domains	SPF configuration and errors
Kondo and colleagues, 2025 [25]	Popular domains, after the 2024 sender requirements	Who had not implemented the required measures
Länge and colleagues, 2026 [30]	The 100,000 most popular domains, daily for eight months	DMARC configuration and its change over time
HTTP Archive Web Almanac, 2025 [50]	Pages from a crawl of the web	Security headers and HTTPS
So and colleagues, 2022 [48]	201 lapsed domains, re-registered	Traffic still sent to them: residual trust
This study	4,406 independent businesses in one city, enumerated from an open places dataset	Email authentication, web transport and headers, announced software, domain lifecycle and lookalikes, on the same businesses, with organisational covariates

Populations are described at the level of detail each study reports; see each reference for its exact sampling.

3. Research questions

1. **RQ1, prevalence.** What share of independent Bakersfield businesses have each basic email and web protection in place?

2. **RQ2, depth.** Among businesses that have begun deploying email authentication, how far do they progress, and where does deployment stall?
3. **RQ3, determinants.** Does protection vary by sector, neighbourhood, email provider, website platform or domain age, and which of these matter once the others are held constant?
4. **RQ4, coherence.** Do protections travel together? Is a business with good email security more likely to have good website security?
5. **RQ5, residual and lookalike risk.** How many local listings point to lapsed domains, and how common are registered lookalikes of local business domains?


4. Methods

4.1 Sampling frame and population

We used the places theme of Overture Maps [41], an open dataset maintained by the Overture Maps Foundation and released under the Community Data License Agreement, Permissive 2.0. It combines contributions from several large mapping and directory sources. We queried every place within a bounding box around metropolitan Bakersfield, then retained places whose address locality is Bakersfield. Each place carries a category hierarchy, an existence confidence score, an operating status, brand attribution where known, and zero or more websites.

The population is defined as *live, independent, private businesses with an address in Bakersfield and their own website domain*. Figure 1 shows how the frame was reduced to it. Chains and franchises were removed in two ways: by brand attribution, and by counting every website domain across all places in California, treating any domain listed by three or more places as a chain wherever its branches are. A business's own domain was required to be a registrable name of exactly two labels; pages on social networks, directories, ordering and booking platforms, and subdomains of other services were removed. Finally, every remaining domain was resolved: domains with no delegation were classed as having no DNS delegation, and domains delegated to known domain parking or resale services as parked. Both were excluded from every measure, as were the few whose name servers could not be determined.

FIGURE 1
How 17,560 listings with a website became 4,406 businesses



Exclusions are applied in the order shown; a listing removed at one step is not counted at later ones. Sample construction follows the reporting style of a participant flow diagram.

Table 2 shows the population by sector, assigned from each place's category hierarchy.

Table 2. Study population by sector

SECTOR	BUSINESSES	SHARE
Professional services	1,383	31.4%
Home services and construction	751	17.0%
Health	524	11.9%
Auto	406	9.2%
Other services	389	8.8%
Food and drink	344	7.8%
Retail	340	7.7%
Beauty, wellness and fitness	269	6.1%

4.2 Measurement

All measurements were taken on 2026-10-18. Each business was measured with passive observation only, using the same information any mail server or web browser obtains in normal operation:

- **DNS.** TXT records at the domain (SPF), at `_dmarc` (DMARC), at `_mta-sts` and `_smtp._tls`, and at `default._bimi`; MX, NS and CAA records; and, for Google Workspace and Microsoft 365 customers, the provider's default DKIM selector (`google`, or `selector1`). SPF records were evaluated recursively to count DNS lookups. Queries went to Cloudflare's public resolvers, falling back to Google's. DNSSEC delegation was read from DS records through Google's DNS over HTTPS service.
- **Registration.** The RDAP record [21] for the domain: creation and expiry dates, status codes (for transfer lock), and registrar.
- **Registry status of domains with no delegation.** Name server delegation was read on 18 October 2026. On 19 October 2026, each domain with no delegation was queried at `https://rdap.org/domain/<name>`, a public redirector that forwards each query to the authoritative registry's RDAP service listed in the IANA bootstrap registry [4]. An HTTP 404 response was recorded as *no registration record returned*; a 200 response whose status values included a hold, redemption period or pending delete as *registered, held*; any other 200 response as *registered*. Rate limits (HTTP 429) and server errors were retried up to five times, waiting four seconds longer after each attempt; each request timed out after 15 seconds; a domain still unanswered after five attempts was recorded as *unknown*. No registration record returned is reported as such, not as confirmed availability.

- **Website.** One GET request to the home page over https, falling back to the www name and then to plain http; one request to the plain http address without following redirects, to see whether it redirects to https; and the TLS handshake on port 443, recording the negotiated protocol, chain validity, expiry, issuer and whether the certificate covers the bare domain. From the response we recorded six security headers, cookie attributes, the platform (from known markers), announced versions of WordPress, PHP and jQuery, third party script hosts and known trackers, mixed content, a link to a privacy policy, and the latest copyright year.
- **Lookalikes.** For each domain, candidates formed by omission, repetition and transposition of characters, common glyph substitutions (for example rn for m, 0 for o), and the same name on eight common endings. Each candidate was resolved for an address record and, if found, mail exchange records. No request was sent to any candidate.

4.3 Operational definitions

A business has **no enforced DMARC policy** when its domain publishes no DMARC policy of quarantine or reject. This measures the owner's published instruction, not whether a forged message would be delivered: receivers may still filter such mail. A business publishes a **reject policy** when it publishes `p=reject` applied to all mail; this is the strongest instruction DMARC offers, not complete protection. The terms *spoofable* and *forgeable* are not used in this paper for that reason. We define **five basic controls**, each a single setting with no licence cost:

1. DMARC enforced: a policy of quarantine or reject.
2. A valid SPF record: present, a single record, within the ten lookup limit, and not ending in `?all` or `+all`.
3. HTTPS enforced: the site serves https and its plain http address redirects to it.
4. A valid certificate: trusted, unexpired and covering the address.
5. No outdated software on show: no WordPress behind the current release branch, no PHP release older than 8.2, and no jQuery older than 3.5.0, where the site announces a version.

Each control is recorded as passed, failed, or not assessed. A website that did not answer is not assessed on the three website controls, since a failed request shows nothing about its certificate or software; email records that no resolver answered for are not assessed on the two email controls. The share passing all five is computed only over businesses where all five could be assessed (4,183 of 4,406); the share failing two or more counts definite failures across all businesses.

4.4 Validity safeguards

- **Absence requires an answer.** A record was counted as missing only when a resolver positively answered that it does not exist. A query that no resolver answered was treated as missing data, never as absence.
- **Dead domains excluded.** Counting a lapsed domain as lacking a DMARC policy would overstate the finding; such domains are reported separately (Section 5.8).
- **Instrument checks.** In a pilot on 182 Kern County businesses, the WordPress version check was found to read the bundled jQuery version from some sites. It was rewritten to read only the generator tag and core files that carry the WordPress version, and every WordPress site was remeasured.

- **Manual review.** The pilot list was read in full to identify national organisations missed by automated filters; the same names are excluded here. Surprising findings, including certificate errors, were inspected by hand before wording was settled.

4.5 Validation against an independent implementation

To estimate the error in the email classifications, a stratified sample was measured again with checkdmarc [54], an open source SPF and DMARC checker with its own DNS client and parser, written independently of this study. It was run with its mail server connection test turned off, so that it too read public DNS only. The sample drew up to 25 businesses at random, with a fixed seed, from each combination of email provider class and DMARC classification, 491 businesses in all, so that rare classes such as reject policies and filtering gateways were checked as well as common ones. The check ran on 2026-10-19.

Table 3. Agreement between the study and an independent implementation, stratified sample

CLASSIFICATION	COMPARED	AGREEMENT	COHEN'S KAPPA
DMARC policy (none, missing, quarantine, reject)	490	99.8% (489)	0.997
SPF record published	490	100.0% (490)	1.000
SPF over the ten lookup limit, or invalid	490	100.0% (490)	1.000

The check found two errors in this study's own instrument, both corrected before the results reported here. First, seven domains publish two DMARC records. The standard treats that as no policy at all, but the study had read the first record, which credited domains with policies they do not effectively have. Second, the SPF lookup counter skipped any included record it had already counted elsewhere in the same evaluation, whereas RFC 7208 counts every evaluation [24]; it undercounted for some domains. All 2,973 stored SPF records were counted again with the corrected rule on 2026-10-19, resolving their included records that day, and 37 counts changed. After both corrections, the one remaining disagreement is a domain whose DMARC and SPF records were both absent when the check ran, a change of record rather than a misclassification.

Agreement between two programs is evidence of correct parsing, not of correctness in every sense, since both read the same DNS. The website measures and the DKIM check were not part of this validation. A review of the same sample by an independent person, using standard lookup tools, has been prepared but not yet completed; its error rates will be published with the next revision.

4.6 Statistical analysis

Proportions are reported with Wilson score 95% confidence intervals [55], which behave well near 0 and 1, where several measures lie. Differences between groups are tested with Pearson's chi-square test of independence without continuity correction, with Cramér's V as the effect size [7]. Groups with fewer than 30 businesses are not reported separately. Two multivariable logistic regression models were fitted by iteratively reweighted least squares, with Wald 95% intervals for odds ratios and McFadden's pseudo R^2 [37]. Model 1 estimates the odds of having no enforced DMARC policy from email provider class, sector, website platform class, log domain age and registrar; Model 2 estimates the odds that a secure website

sends HSTS from DMARC enforcement, platform class and sector. Pairwise associations between seven controls are summarised as odds ratios with a 0.5 continuity correction. Monotonic association between domain age and controls passed uses Spearman's rank correlation.

The analysis was implemented directly on numpy rather than with SciPy or statsmodels, because the study machine's application control policy blocks those libraries' compiled extensions. The chi-square tail uses the regularised incomplete gamma function [44]. The implementation was checked against known critical values and against a simulated logistic model with known coefficients before use; both matched to three decimal places. The models have not yet been rerun in a standard statistical package on another machine; that rerun is planned and its results will be published with the next revision.

4.7 Ethics


The study follows the principles of the Menlo Report [11]: respect for persons, beneficence, justice and respect for law and public interest. Every observation is of information each business publishes to the world for machines to read, collected in the same manner and volume as a single visitor or mail server. Nothing was scanned, no vulnerability was tested, no login was attempted, no form was submitted, and no personal data was collected. Results are reported only in aggregate; no business is named here or elsewhere, and per business data is kept privately and will not be published or sold. A business that believes it was included may request its own results.

5. Results

5.1 Overall posture

KEY POINTS

- Only **357 of 4,183** businesses (8.5%) pass all five basic controls.
- The usual profile is **three or four out of five**: the website passes and the email fails.
- An enforced DMARC policy is by far the least common control, held by 13.9% of businesses.

FIGURE 2**How many of the five basic controls each business passes**

Distribution of controls passed across the 4,183 businesses where all five could be assessed. The orange column marks the businesses passing all five.

FIGURE 3**Share of businesses passing each basic control**

Bars show the share passing; whiskers show the Wilson 95% confidence interval.


5.2 Email authentication (RQ1, RQ2)

KEY POINTS

- **86.1% (95% CI 85.0% to 87.1%)** of business domains lack an enforced DMARC policy.
- Deployment stalls at every stage: of 4,402 domains, 2,837 publish SPF, 1,396 publish DMARC, 612 enforce it, and only 201 publish a reject policy *and* receive the reports that make the policy manageable.
- **61.4%** of monitor only DMARC records send their reports nowhere: the step was started and abandoned.
- A DKIM key at the default selector is found far less often for Google Workspace customers (24.1%) than for Microsoft 365 customers (39.6%).

FIGURE 4

Where email authentication deployment stops




Each stage as a share of all businesses with readable email records. Hover a bar for its share of the previous stage.

The funnel in Figure 4 is the central descriptive result of the study. Roughly a third of businesses (35.6%) publish no SPF record at all. Of those that do, about half publish DMARC. Of those, fewer than half enforce

it, and of those, fewer than half reject. The loss at each step is consistent with a deployment path that is begun, often by a provider's setup wizard or a deliverability prompt, and then not completed.

FIGURE 5

DMARC policy published




Share of 4,402 businesses with readable email records. Policies are ordered from no policy to the strictest. One record with an unrecognised policy, which receivers treat as no policy, is counted with the missing.

SPF, where present, is usually syntactically sound but lenient. Table 4 classifies every SPF record. The ten lookup limit is rarely breached (50 records), but those records fail permanently, and their owners almost certainly believe they are protected. The median record requires 2 lookups (Figure 6).

Table 4. SPF status among 4,402 businesses

STATUS	MEANING	BUSINESSES	SHARE (95% CI)
Missing	No SPF record published	1,565	35.5% (34.2% to 37.0%)
Soft fail (~all)	Unlisted senders marked suspicious	1,373	31.2% (29.8% to 32.6%)
Strict (-all)	Unlisted senders fail	1,264	28.7% (27.4% to 30.1%)
Open (?all or +all)	Ending that protects nothing	110	2.5% (2.1% to 3.0%)
Broken	Two records, or over ten lookups: permanent error	90	2.0% (1.7% to 2.5%)

FIGURE 6

DNS lookups required to evaluate each SPF record

Records needing more than ten lookups (orange) produce a permanent error under RFC 7208 and protect nothing. The last column groups twelve or more.

Both Google Workspace and Microsoft 365 sign mail with the customer's own domain only after an administrator generates a key, publishes it in DNS and switches signing on. The two populations nonetheless differ markedly: a DKIM key is published at the provider's default selector for 24.1% (95% CI 21.2% to 27.2%) of Google Workspace customers and 39.6% (95% CI 36.6% to 42.7%) of Microsoft 365 customers ($\chi^2 = 47.86$, $p < 0.001$). These are lower bounds on key publication, not measures of signing: a business using a custom selector would be missed, and a published key does not prove that outgoing mail is signed with it [52].

Finally, 1,047 businesses receive no mail at their website domain, typically because they use a personal mailbox. A domain that sends no mail can be locked with two records (`v=spf1 -all` and a reject policy). 98.8% (95% CI 97.9% to 99.3%) of these domains are not locked, so receivers have no instruction from the owner to reject mail claiming to come from them, even though the owner never uses them for email. Newer transport protections are almost absent: MTA-STS is published by 10 of 3,355 mail receiving domains, and BIMI by 9.

5.3 Email provider (RQ3)

KEY POINTS

- Provider is the strongest single correlate of protection in the study (Cramér's $V = 0.263$, the largest of any grouping).
- Businesses behind a **filtering gateway** (63.8%) or on **Microsoft 365** (71.5%) lack an enforced policy far less often than those on **Google Workspace** (91.8%) or their web host's mail (90.8%).

FIGURE 7**Share of business domains with no enforced DMARC policy, by email provider**

Provider inferred from MX records. Filtering gateways are Proofpoint, Mimecast, Barracuda and AppRiver. $\chi^2 = 305.01$, $df = 5$, $p < 0.001$, $V = 0.263$.


The detail matters. Within filtering gateways, Proofpoint and Mimecast customers are the best protected businesses in the study, while AppRiver, often resold with hosted mailboxes, looks like the web host group (Appendix B). We read provider as a proxy for *whether a professional configured the domain*: gateways and Microsoft 365 tenants are far more often deployed by an IT provider than a Google Workspace account opened by an owner, or the mailbox bundled with a website plan.

5.4 Websites: a transport baseline by default (RQ1, RQ3)


KEY POINTS

- Encryption is nearly universal: **93.4%** (95% CI 92.6% to 94.2%) negotiate TLS 1.3.


- But **25.3%** (95% CI 24.0% to 26.6%) do not send plain http visitors to the secure site, and **45.3%** send none of six standard security headers.
- Hosted site builders send at least one security header **87.7%** of the time; self managed sites, **40.1%**. Custom sites redirect to https only **56.1%** of the time.

FIGURE 8**Security headers sent by business websites**

Share of reachable websites sending each header (HSTS: share of https sites). Whiskers are 95% intervals.

FIGURE 9**Number of the six security headers each website sends**

Distribution across 4,186 reachable websites. The median is 1.

FIGURE 10**The platform paradox: who protects the website, and who protects the email****Website sends a security header****No enforced DMARC policy**

Upper panel: share of reachable websites sending at least one security header. Lower panel: share of businesses with no enforced DMARC policy. Hosted builders lead on the first and trail on the second.

Figure 10 is the clearest evidence for this paper's title. Wix, Squarespace, Shopify and similar builders apply transport security and response headers to every customer site automatically; their customers need do nothing, and they score best on the web. Yet those same businesses are the most likely to lack an enforced DMARC policy (91.8%), because no builder configures a customer's email domain. Self managed sites, usually WordPress, are the reverse: worst on headers, best on email (79.7% without an enforced policy). A plausible reading is that a self managed site implies a web professional, an agency or IT contractor, who also has access to the domain's DNS and is more likely to have set up mail authentication. The difference survives adjustment for provider, sector and domain age (Model 1, OR 0.46, $p < 0.001$).

89.9% of reachable websites present a valid, trusted certificate. The failures are visible ones: 320 sites serve no HTTPS at all, and 200 (5.0%) present a certificate warning on a common form of their address. Most of these cover the www name but not the bare domain, so a customer who omits www sees a full page browser warning; most of the remainder have expired. Let's Encrypt issues 55.9% of all certificates observed. With maximum lifetimes falling to 100 days in March 2027 and 47 days by March 2029 [6], certificates renewed by hand will lapse more often unless renewal is automated.


5.5 Software exposure

KEY POINTS

- 33.7% of websites publicly announce a software version number.
- 30.1% (95% CI 27.1% to 33.2%) of WordPress sites showing a version are behind the current release; 80 are more than five releases behind.
- Of sites announcing PHP, 47.0% run a release with no security support.

FIGURE 11

How far behind the current WordPress release sites are



WordPress sites that announce their version (n = 878). Releases counted as minor versions behind 7.1.3. Orange columns are more than five releases behind.

The majority of WordPress sites in the city are current, which reflects WordPress's automatic minor updates. The tail is long: the oldest release observed dates from 2013. Among sites announcing a jQuery version, 18.5% (247) load a release before 3.5.0. Platform shares are shown in Appendix C.


5.6 Sector and neighbourhood (RQ3)

KEY POINTS


- Sector differences are real but modest ($V = 0.125$): from **80.7%** without an enforced policy in professional services to **93.9%** in food and drink.
- **Nearly 9 in 10 health care practices** (87.8%) lack an enforced DMARC policy.
- Neighbourhood barely matters ($V = 0.085$): every ZIP code lies between 81.5% and 91.6%.

FIGURE 12

Share of business domains with no enforced DMARC policy, by sector



$\chi^2 = 68.79$, $df = 7$, $p < 0.001$, Cramér's $V = 0.125$. Whiskers are 95% intervals.

FIGURE 13**Share of business domains with no enforced DMARC policy, by ZIP code**

$\chi^2 = 31.42$, $df = 10$, $p < 0.001$, Cramér's $V = 0.085$. Statistically detectable, but small: the problem is citywide.

Table 5. Key measures by sector


SECTOR	N	NO ENFORCED DMARC	ALL FIVE BASICS	ANY SECURITY HEADER	WORDPRESS BEHIND
Food and drink	343	93.9%	1.8%	62.8%	45.0% (n = 40)
Auto	406	90.9%	5.2%	57.1%	28.8% (n = 73)
Beauty, wellness and fitness	268	90.7%	4.7%	55.6%	34.9% (n = 43)
Retail	339	88.5%	6.9%	57.9%	32.6% (n = 46)
Other services	389	88.2%	7.3%	53.3%	20.3% (n = 74)
Health	524	87.8%	6.8%	54.5%	35.3% (n = 119)
Home services and construction	751	85.0%	9.2%	53.4%	29.8% (n = 178)
Professional services	1,382	80.7%	13.1%	52.2%	27.9% (n = 305)

5.7 What predicts protection: a multivariable model (R03)

KEY POINTS

- Holding sector, platform, domain age and registrar constant, **Microsoft 365** customers have 0.27 times the odds of lacking an enforced policy that Google Workspace customers have, and **filtering gateway** customers 0.28 times.
- Auto (2.05), food and drink (1.75) and health (1.44) businesses carry higher odds than professional services, independent of provider.
- Older domains are better protected; the registrar makes no difference.

FIGURE 14

Model 1: odds that a business domain has no enforced DMARC policy

0.77 to 1.15

0.1 0.25 0.5 1 2 4

Filled markers are significant at $p < 0.05$; grey markers are not. Blue lowers the odds, orange raises them.

Multivariable logistic regression, $n = 3,448$ businesses with readable email and registration records (2,949 without an enforced policy). McFadden pseudo $R^2 = 0.1004$; likelihood ratio $\chi^2 = 286.24$, $p < 0.001$. Odds ratios below 1 mean lower odds of lacking enforcement.

The model explains a modest share of variation (pseudo R^2 0.1004), as expected for an outcome that four in five businesses share, but its pattern is clear. Provider class dominates. Website platform remains significant after provider is controlled, consistent with the reading that a self managed site signals professional involvement. Domain age has a protective effect (OR 0.77 per unit of log years, $p = 0.002$), though Section 5.8 shows the relationship is not simple. Registration at GoDaddy, which holds half of the city's business domains, has no detectable effect (OR 0.94, $p = 0.561$): the registrar is where the domain lives, not who configures it.

5.8 Domain lifecycle, lapsed listings and lookalikes (RQ5)

KEY POINTS

- **1,117 of 5,722 listings** point to a domain with no DNS delegation, and 136 more to a parked one.
- Protection follows a **U shaped curve** with domain age: newest and oldest domains do best, those 2 to 10 years old worst.
- **25.9% (95% CI 24.6% to 27.2%)** of businesses have a registered lookalike domain configured for email.

FIGURE 15**Share of business domains with no enforced DMARC policy, by domain age**

Age from the RDAP registration date. $\chi^2 = 158.96$, $df = 4$, $p < 0.001$, $V = 0.214$. Spearman correlation between age and basic controls passed: $\rho = 0.111$ ($p < 0.001$).

The U shape in Figure 15 has two plausible explanations, which this cross sectional study cannot separate. The oldest domains, over twenty years, belong disproportionately to established firms more likely to retain an IT provider. The newest, under two years, were set up after Google and Yahoo began requiring DMARC of bulk senders in February 2024 [16], when setup guides, provider wizards and deliverability warnings started prompting new domain owners to publish a record. Domains in the middle were configured before that shift and never revisited. A repeat of this study will show whether the newest cohort keeps its advantage.

The lapsed listings are a distinct risk. Of 5,722 independent business listings with their own domain, 1,117 (19.5%) name a domain with no delegation, and 136 name one held by a domain seller. Some of these businesses have closed; others moved to a new domain and never updated their listings. Where the registry has released the name, anyone can register it and would then receive mail sent to the old addresses and control the website that maps, directories and old printed material still point to. This is the residual trust problem described by Lever and colleagues and quantified by So and colleagues [32, 48], observed here at the scale of one city's local listings. No delegation does not by itself mean a domain is available, since a registered domain can sit on hold or in a grace period without name servers, so each of these domains was looked up in the registry through RDAP [21]: 995 returned no registration record, 70 were registered but on hold, in redemption or pending deletion, 29 were registered without such a status, and 23 could not be determined (Section 4.2 gives the protocol). No registration record is not the same as confirmed availability, since registries can reserve, block or restrict names, but most such names can ordinarily be registered by anyone.

Domains in the study are old (median 12.3 years) and nearly always transfer locked (93.9%). DNSSEC (3.4%) and CAA (1.1%) are rare. Registrar concentration is high: GoDaddy holds 1,754 of 3,485 readable registrations (50.3%), Registrar concentration is not account concentration: each business has its own account. It does mean that the protection of half the city's business domains rests on one provider's login

protections, recovery and support processes, and that for each of those businesses one account controls its website and every email record.

Across 269,689 lookalike candidates, 3,889 were live and 2,426 had mail records. 34.1% of businesses have at least one registered typo domain and 25.9% one configured for mail. Many of these are benign, belonging to unrelated businesses with similar names or to the business itself, so this is a measure of exposure, not of abuse.


5.9 Do protections travel together? (RQ4)

KEY POINTS

- Email controls cluster: businesses enforcing DMARC are **12 times** as likely in odds to publish a DKIM key, and 4.4 times to have strict SPF.
- Web controls cluster: HTTPS redirect with HSTS (OR 6.3); HSTS with Content Security Policy (OR 5.6).
- **Across the two families, a weak link that platform hides:** unadjusted, DMARC enforcement with HSTS has OR 0.91 and with Content Security Policy OR 0.85, neither significant; adjusted for platform and sector, DMARC enforcement with HSTS has OR 1.43.

FIGURE 16

Pairwise association between seven controls



Each cell is the odds ratio for having both controls. CSP is Content Security Policy. Blue cells are controls that travel together, orange cells controls that avoid each other, and dashed cells are not significant at $p < 0.05$. Values above 3 are bold.

Figure 16 shows two blocks. The upper left block, the email controls, is strongly blue: a business that enforces DMARC publishes a default selector DKIM key 77.3% of the time, against 22.0% for one that does not. The lower right block, the web controls, is also blue. The rectangle between them is mostly grey: unadjusted, knowing that a business has protected its email tells us little about whether it has protected its website.

Model 2 refines this. Hosted builders send HSTS by default and rarely enforce DMARC, which pulls the raw association toward zero. Holding platform and sector constant, DMARC enforcement is associated with modestly higher odds of HSTS (OR 1.43, 95% CI 1.16 to 1.76, $p < 0.001$). This masking is the more interesting result. In the raw data, platform defaults hide a positive association between the two families: within the same kind of platform, businesses that enforce DMARC are somewhat more likely to send HSTS, consistent with some shared attention to security or shared professional help. The adjusted association is still several times weaker than the associations within each family. One explanation, which this study cannot test, is that website and email settings are often set by different parties, or one of them by nobody; establishing who is responsible requires asking owners and providers directly.


5.10 Privacy signals

KEY POINTS

- 35.5% of websites load at least one tracking or advertising script, most often Google Analytics.
- Of those, **44.1% (95% CI 41.6% to 46.7%)** have no privacy policy link on the home page.

FIGURE 17

Tracking and advertising scripts found on business home pages



Websites loading each script, of 4,186 reachable sites. A site may load several.

California's consumer privacy law applies chiefly to larger businesses, and most businesses in this population fall below its thresholds. The finding matters for a different reason: the terms of the major analytics and advertising platforms require a privacy disclosure, and its absence is one of the first things an insurer or a larger client reviews. 41.5% of sites that set cookies over https set at least one without the Secure attribute, and 11.7% of https pages load some content over plain http.

6. Discussion

6.1 Secure by default, exposed by omission

The web and email followed different histories. Over the last decade, the web acquired secure defaults: free automated certificates [1], browsers that warn on plain http, hosting platforms that apply encryption and headers to every customer. A small business that does nothing still receives most of the benefit, and the data shows it: 93.4% TLS 1.3, and the best header coverage on the platforms where owners do the least. Email acquired no equivalent. In general, registrars, mailbox providers and website builders do not publish an enforced DMARC policy for a customer's domain by default. The owner must act, or must have someone act for them, and 86.1% have not. The pressures described in Section 2.1 largely pass these businesses by. The mailbox provider requirements apply to senders of more than 5,000 messages a day, a volume few local businesses approach, and are satisfied by a monitoring policy, which asks receivers not to change how they handle failing mail; that is consistent with monitoring being the most common configuration among popular domains [30] and with the 783 monitor only records found here, most of them with no reporting address. Card payment rules and government guidance point toward reject [9, 43], but reach an owner only through a compliance questionnaire or an IT provider.

This reframes the problem. The common explanation for poor small business security is that owners do not care or do not know. The patterns here point elsewhere: protection follows defaults where defaults exist, and, on the evidence of provider and platform type, appears to follow professional involvement where they do not; that last step is a hypothesis this study did not test directly. Owners are behaving as Herley's rational user [18] would, attending to what is visible (a padlock, a working website) and leaving alone what is not (a DNS record no customer sees). Email forgery is invisible to the business until it is used against them.

6.2 How Bakersfield compares with the wider internet

Table 6 sets this study's measures beside the most comparable recent figures for popular domains and web pages. Populations, dates and definitions differ, so the comparison is indicative only. Its pattern is consistent with the rest of this paper: on the website, Bakersfield's small businesses look much like the web at large; on email, they publish DMARC at about half the rate of popular domains and publish a reject policy with reporting at about a third of the rate.

Table 6. Bakersfield small businesses against recent large scale measurements

MEASURE	BAKERSFIELD	COMPARISON
DMARC record published	31.7%	58.4% of the 100,000 most popular domains, end of an eight month scan from April 2025 [30]
Reject policy with reporting	4.6%	13.1% to 14.8% of the same domains [30]
SPF record published	64.4%	56.5% of 12 million domains, 2023 [10]
SPF broken	2.0%	2.9% of SPF records with errors, 2023; definitions differ [10]
Content Security Policy	21.7%	21.9% of pages, 2025 [50]
HSTS	41.3% of https sites	36% of mobile pages, 2025 [50]

6.3 Weakly linked: email and website controls

The co-occurrence structure (Figure 16) and Model 2 suggest that a small business does not have one security posture but at least two, which we hypothesise are often set by different parties. This has a practical consequence for anyone assessing such a business: checking the website, which is what most free online scanners and most owners look at, says little about the email domain. It also suggests, though this study did not measure it, where responsibility may fall through. The web designer considers email out of scope; the mailbox provider considers DNS the owner's; the registrar hosts the records but configures none of them.

6.4 Who configures matters more than what was bought

Google Workspace and Microsoft 365 offer equivalent authentication capabilities, yet their customers differ sharply in outcome, and the difference holds after accounting for sector, platform and age. Both require the same deliberate steps to sign mail and to publish DMARC, so the most economical explanation is the setup path rather than the product: Microsoft 365 tenants for small businesses are commonly deployed by IT providers, while Google Workspace is more often opened by the owner. The best protected group of all, businesses behind enterprise filtering gateways, is a product almost always installed and maintained by an IT provider. Provider type is not measured involvement, and this interpretation should be tested directly. For local IT firms, this is both an indictment and an opportunity: their clients are measurably better protected, and the remaining gap is a few records per client.

6.5 Sector

Professional services, which handle money and confidential documents, are best protected; food and drink, beauty and auto services least. The health result deserves emphasis: 87.8% of medical, dental and veterinary practices publish no enforced DMARC policy, so receivers have no instruction from the practice to reject mail that fakes its domain, mail patients have every reason to trust when it concerns a bill or an appointment.

6.6 Lapsed domains

One in five business listings pointing to a domain with no delegation is, to our knowledge, a previously unreported measure at the local level. It is partly a data quality finding about map listings, but it is also a security finding: each such domain is a name that customers and suppliers still associate with a real business, and where the registry has released it, anyone may register it.

7. Threats to validity

- **Coverage.** The frame includes only businesses listed in Overture Maps with a website. Businesses without a website, likely less digitally engaged, are absent; if anything this biases measured protection upward.
- **Chain detection.** A chain with one or two California listings and no brand attribution could remain in the population. Manual review of the pilot and the three listing threshold limit this, but cannot exclude it.
- **What the DMARC measure means.** This study measures published configuration, not attacks, deliveries or losses. Receivers apply their own filtering, so forged mail from a domain without an enforced policy may still be blocked, and a reject policy does not stop compromised accounts, lookalike domains or display name impersonation. The measure captures what the domain owner has instructed, not what any receiver does.
- **DKIM measurement.** DKIM is measured only at default selectors for two providers; a business using a custom selector is recorded as having no DKIM key found, and a key that is found does not show that mail is signed. DKIM figures are therefore lower bounds on key publication.
- **Version detection.** Software versions are counted only where announced. Sites that hide versions are neither credited nor penalised, and announced versions can lag the actual software.
- **Validation scope.** The email classifications were checked against an independent implementation on a stratified sample (Section 4.5); website measures, the DKIM check and the registry status protocol were not, and the human review of the sample is not yet complete.
- **Plugins.** WordPress plugins, where most vulnerabilities in that ecosystem are found [42], were not assessed; software exposure is therefore a lower bound.
- **Provider inference.** Provider is inferred from MX records. A business using a gateway in front of Microsoft 365 is counted under the gateway.

- **Selection, and what the intervals mean.** The population is businesses that appear in Overture Maps with their own website, after the exclusions in Figure 1; it is not every business in Bakersfield. The chain rule (a domain at three or more California locations) can exclude locally owned businesses with several branches. Confidence intervals describe sampling error within this population only; they do not correct for how the population was selected, and a large sample is not by itself representative of the city.
- **Lapsed domains.** Registry status was read once through RDAP and can change daily; a domain released after the check, or registered by someone else since, would be misclassified.
- **Causality.** All associations are cross sectional. Provider and platform are interpreted as proxies for professional involvement; that interpretation is plausible but not directly measured.
- **Single day.** All measures reflect 2026-10-18. DNS and websites change; individual results may differ on another day.
- **Generalisation.** Bakersfield is a mid sized California city with a large professional services and agricultural economy. The pilot across Kern County produced results within a few points, but other regions may differ.

8. Recommendations

8.1 For business owners

1. **Publish DMARC, then enforce it.** Start at p=none with a reporting address; after a few weeks of clean reports move to quarantine, then reject.
2. **Turn on DKIM.** Google Workspace: Admin console, Apps, Google Workspace, Gmail, Authenticate email. Microsoft 365: Microsoft Defender portal, Email authentication settings, DKIM.
3. **Correct SPF:** one record, every sending service listed, within ten lookups.
4. **Lock domains that send no mail,** including old domains: v=spf1 -all and a reject policy.
5. **If you take card payments,** ask your payment processor which PCI DSS self assessment applies to you; the current standard expects anti-phishing controls and names DMARC, SPF and DKIM as examples [43].
6. **Protect the registrar login** with two step verification, and keep auto renewal and a current card on file.
7. **Check every map and directory listing** points to a domain you control today; renew old domains rather than letting them lapse.
8. **Redirect http to https,** cover the address with and without www, enable HSTS, and automate certificate renewal before lifetimes fall to 100 days in March 2027 [6].
9. **Update or retire self managed software,** or move to a host that updates it automatically.

8.2 For IT providers and web designers

Treat the domain, not the website or the mailbox, as the unit of responsibility. Every client onboarding should end with SPF, DKIM and an enforced DMARC policy, and every website handover should state who

owns the domain's mail records. The gap between the two silos is where client fraud happens.

8.3 For platforms and registrars

Website builders already apply secure defaults to the web. Registrars and builders that host DNS could publish a monitoring DMARC record by default, and a null SPF and reject policy for domains with no mail configured, at no cost to the customer. Google Workspace could enable DKIM by default for verified domains. Each of these would change the published configuration of thousands of businesses in this one city without any owner action. This study suggests, but did not test, that such defaults would reach more of them than advice to owners does.

9. Conclusion

In a population of 4,406 independent businesses in one California city, the website usually meets a basic transport baseline and the email domain usually lacks an enforced authentication policy. 86.1% of business domains publish no DMARC policy of quarantine or reject; few pass five free, basic controls. The pattern is not random. Website controls follow platform defaults; email authentication appears to follow professional involvement, a hypothesis that needs direct measurement; and the two are only weakly linked once platform is taken into account. No intervention was tested here; the patterns suggest, as hypotheses for that work, that defaults from platforms and completeness from the professionals who already serve these businesses may reach more of them than advice to owners. This study will be repeated annually so that change, or its absence, can be measured.

Statements

Competing interests. The author is the owner of Mitivanta, a firm that sells cybersecurity assessments to small businesses, including in Bakersfield. The study design, measures and thresholds were fixed before results were known, the method is published in full, and no business is identified.

Data availability. Aggregate results, the figures' underlying data and the full model outputs are published on this page. The measurement and analysis code is held by Mitivanta and available to researchers on request. Per business data will not be released, to protect the businesses measured.

Acknowledgements. Business listings are from the Overture Maps Foundation, used under CDLA Permissive 2.0. Registration data is from the public RDAP service.

Check your own business. The email, website and lookalike checks used in this study are available free at mitivanta.com/tools.

Suggested citation. Wattenbarger, S. (2026). The Bakersfield Cybersecurity Report 2026: Secure by Default, Exposed by Omission. Mitivanta Research. <https://mitivanta.com/bakersfield-cybersecurity-report>

References

1. Aas, J., Barnes, R., Case, B., Durumeric, Z., Eckersley, P., Flores-López, A., Halderman, J. A., Hoffman-Andrews, J., Kasten, J., Rescorla, E., Schoen, S., and Warren, B. (2019). Let's Encrypt: An Automated Certificate Authority to Encrypt the Entire Web. Proceedings of the ACM Conference on Computer and Communications Security (CCS).
2. Agten, P., Joosen, W., Piessens, F., and Nikiforakis, N. (2015). Seven Months' Worth of Mistakes: A Longitudinal Study of Typosquatting Abuse. Proceedings of the Network and Distributed System Security Symposium (NDSS).
3. Ashiq, M. I., Li, W., Fiebig, T., and Chung, T. (2023). You've Got Report: Measurement and Security Implications of DMARC Reporting. Proceedings of the 32nd USENIX Security Symposium, 4123 to 4137. www.usenix.org/conference/usenixsecurity23/presentation/ashiq
4. Blanchet, M. (2022). Finding the Authoritative Registration Data Access Protocol (RDAP) Service. RFC 9224. Internet Engineering Task Force. www.rfc-editor.org/info/rfc9224
5. Brotman, A. (Ed.) (2026, May). DMARC Aggregate Reporting. RFC 9990. Internet Engineering Task Force. www.rfc-editor.org/info/rfc9990
6. CA/Browser Forum (2025, April 11). Ballot SC081v3: Introduce Schedule of Reducing Validity and Data Reuse Periods. cabforum.org/2025/04/11/ballot-sc081v3-introduce-schedule-of-reducing-validity-and-data-reuse-periods/
7. Cramér, H. (1946). Mathematical Methods of Statistics. Princeton University Press.
8. Crocker, D., Hansen, T., and Kucherawy, M. (2011). DomainKeys Identified Mail (DKIM) Signatures. RFC 6376. Internet Engineering Task Force. www.rfc-editor.org/info/rfc6376
9. Cybersecurity and Infrastructure Security Agency, National Security Agency, Federal Bureau of Investigation, and Multi-State Information Sharing and Analysis Center (2023, October 18). Phishing Guidance: Stopping the Attack Cycle at Phase One. www.cisa.gov/resources-tools/resources/phishing-guidance-stopping-attack-cycle-phase-one
10. Czybik, S., Horlboge, M., and Rieck, K. (2023). Lazy Gatekeepers: A Large-Scale Study on SPF Configuration in the Wild. Proceedings of the ACM Internet Measurement Conference (IMC), 344 to 355. doi.org/10.1145/3618257.3624827
11. Dittrich, D., and Kenneally, E. (2012). The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. U.S. Department of Homeland Security.
12. Durumeric, Z., Adrian, D., Mirian, A., Kasten, J., Bursztein, E., Lidzborski, N., Thomas, K., Eranti, V., Bailey, M., and Halderman, J. A. (2015). Neither Snow Nor Rain Nor MITM: An Empirical Analysis of Email Delivery Security. Proceedings of the ACM Internet Measurement Conference (IMC).
13. Federal Bureau of Investigation, Internet Crime Complaint Center (2024, September 11). Business Email Compromise: The \$55 Billion Scam. Public Service Announcement I-091124-PSA. www.ic3.gov/PSA/2024/PSA240911

14. Federal Bureau of Investigation, Internet Crime Complaint Center (2026). 2025 IC3 Annual Report. www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf
15. Foster, I. D., Larson, J., Masich, M., Snoeren, A. C., Savage, S., and Levchenko, K. (2015). Security by Any Other Name: On the Effectiveness of Provider Based Email Security. Proceedings of the ACM Conference on Computer and Communications Security (CCS).
16. Google. Email sender guidelines. Google Workspace Admin Help. Requirements in force from February 1, 2024. support.google.com/a/answer/81126
17. Hallam-Baker, P., Stradling, R., and Hoffman-Andrews, J. (2019). DNS Certification Authority Authorization (CAA) Resource Record. RFC 8659. Internet Engineering Task Force. www.rfc-editor.org/info/rfc8659
18. Herley, C. (2009). So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users. Proceedings of the New Security Paradigms Workshop (NSPW).
19. Herr, T., and Levine, J. (Eds.) (2026, May). Domain-based Message Authentication, Reporting, and Conformance (DMARC). RFC 9989. Internet Engineering Task Force. Obsoletes RFC 7489. www.rfc-editor.org/info/rfc9989
20. Hodges, J., Jackson, C., and Barth, A. (2012). HTTP Strict Transport Security (HSTS). RFC 6797. Internet Engineering Task Force. www.rfc-editor.org/info/rfc6797
21. Hollenbeck, S., and Newton, A. (2021). JSON Responses for the Registration Data Access Protocol (RDAP). RFC 9083. Internet Engineering Task Force. www.rfc-editor.org/info/rfc9083
22. Hu, H., and Wang, G. (2018). End to End Measurements of Email Spoofing Attacks. Proceedings of the 27th USENIX Security Symposium.
23. Hureau, O., Bayer, J., Duda, A., and Korczyński, M. (2024). Spoofed Emails: An Analysis of the Issues Hindering a Larger Deployment of DMARC. Passive and Active Measurement (PAM 2024). Lecture Notes in Computer Science, 232 to 261. doi.org/10.1007/978-3-031-56249-5_10
24. Kitterman, S. (2014). Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1. RFC 7208. Internet Engineering Task Force. www.rfc-editor.org/info/rfc7208
25. Kondo, D., Shibuya, Y., Yamaguchi, R. S., Ishihara, T., Sekiya, Y., Nakata, T., and Asami, T. (2025). Who Did Not Implement Email Security Measures After Google's New Email Sender Guidelines?: A Large-Scale Measurement Study. Network Traffic Measurement and Analysis Conference (TMA). doi.org/10.23919/TMA66427.2025.11097004
26. Kranch, M., and Bonneau, J. (2015). Upgrading HTTPS in Mid Air: An Empirical Study of Strict Transport Security and Key Pinning. Proceedings of the Network and Distributed System Security Symposium (NDSS).
27. Krombholz, K., Mayer, W., Schmiedecker, M., and Weippl, E. (2017). "I Have No Idea What I'm Doing": On the Usability of Deploying HTTPS. Proceedings of the 26th USENIX Security Symposium.
28. Kucherawy, M., and Zwicky, E. (2015). Domain-based Message Authentication, Reporting, and Conformance (DMARC). RFC 7489. Internet Engineering Task Force. Obsoleted by RFC 9989. www.rfc-editor.org/info/rfc7489

29. Kumaran, N. (2023, October 3). New Gmail protections for a safer, less spammy inbox. Google. blog.google/products/gmail/gmail-security-authentication-spam-protection/
30. Länge, T., Ballreich, F. L., Hennig, A., Mayer, P., and Volkamer, M. (2026). From Matrix to Metrics: Introducing and Applying a Configuration Matrix to Evaluate DMARC Policies. NDSS Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb). dx.doi.org/10.14722/madweb.2026.23031
31. Lauinger, T., Chaabane, A., Arshad, S., Robertson, W., Wilson, C., and Kirda, E. (2017). Thou Shalt Not Depend on Me: Analysing the Use of Outdated JavaScript Libraries on the Web. Proceedings of the Network and Distributed System Security Symposium (NDSS).
32. Lever, C., Walls, R., Nadj, Y., Dagon, D., McDaniel, P., and Antonakakis, M. (2016). Domain-Z: 28 Registrations Later. Measuring the Exploitation of Residual Trust in Domains. Proceedings of the IEEE Symposium on Security and Privacy.
33. Lim, K., Kwon, Y., and Kim, D. (2023). A Longitudinal Study of Vulnerable Client-side Resources and Web Developers' Updating Behaviors. Proceedings of the ACM Internet Measurement Conference (IMC), 162 to 180. doi.org/10.1145/3618257.3624804
34. Lim, K., Sommese, R., Jonker, M., Mok, R., claffy, k., and Kim, D. (2025). Registration, Detection, and Deregistration: Analyzing DNS Abuse for Phishing Attacks. arXiv:2502.09549 (preprint). arxiv.org/abs/2502.09549
35. Margolis, D., Brotman, A., Ramakrishnan, B., Jones, J., and Risher, M. (2018). SMTP TLS Reporting. RFC 8460. Internet Engineering Task Force. www.rfc-editor.org/info/rfc8460
36. Margolis, D., Risher, M., Ramakrishnan, B., Brotman, A., and Jones, J. (2018). SMTP MTA Strict Transport Security (MTA-STS). RFC 8461. Internet Engineering Task Force. www.rfc-editor.org/info/rfc8461
37. McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*. Academic Press.
38. MITRE. CVE-2020-11022 and CVE-2020-11023 (jQuery cross site scripting). Common Vulnerabilities and Exposures.
39. National Institute of Standards and Technology (2024, February 26). The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29. doi.org/10.6028/NIST.CSWP.29
40. National Institute of Standards and Technology (2024, February 26). NIST Cybersecurity Framework 2.0: Small Business Quick-Start Guide. NIST SP 1300. doi.org/10.6028/NIST.SP.1300
41. Overture Maps Foundation. Places theme, release 2026-09-23.1. Community Data License Agreement, Permissive 2.0. overturemaps.org
42. Patchstack (2026). State of WordPress Security in 2026. patchstack.com/whitepaper/state-of-wordpress-security-in-2026/
43. PCI Security Standards Council (2024, June). Payment Card Industry Data Security Standard: Requirements and Testing Procedures, Version 4.0.1. Requirement 5.4.1. www.pcisecuritystandards.org/document_library/
44. Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes: The Art of Scientific Computing* (3rd ed.). Cambridge University Press.

45. Puneeth (2025, April 2; updated April 29). Strengthening Email Ecosystem: Outlook's New Requirements for High-Volume Senders. Microsoft Defender for Office 365 Blog. techcommunity.microsoft.com/blog/microsoftdefenderforoffice365blog/strengthening-email-ecosystem-outlook%E2%80%99s-new-requirements-for-high%E2%80%90volume-senders/4399730
46. Rescorla, E. (2018). The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446. Internet Engineering Task Force. www.rfc-editor.org/info/rfc8446
47. Rose, S., Nightingale, S., Garfinkel, S., and Chandramouli, R. (2019, February). Trustworthy Email. NIST Special Publication 800-177 Revision 1. National Institute of Standards and Technology. doi.org/10.6028/NIST.SP.800-177r1
48. So, J., Miramirkhani, N., Ferdman, M., and Nikiforakis, N. (2022). Domains Do Change Their Spots: Quantifying Potential Abuse of Residual Trust. Proceedings of the IEEE Symposium on Security and Privacy, 2130 to 2144. doi.org/10.1109/SP46214.2022.9833609
49. Szurdi, J., Kocso, B., Cseh, G., Spring, J., Felegyhazi, M., and Kanich, C. (2014). The Long "Taile" of Typosquatting Domain Names. Proceedings of the 23rd USENIX Security Symposium.
50. Vanderlinden, V., and Franken, G. (2026, January 15). Security. In HTTP Archive, Web Almanac 2025, chapter 9. almanac.httparchive.org/en/2025/security
51. Verizon (2026). 2026 Data Breach Investigations Report. www.verizon.com/business/resources/reports/dbir/
52. Wang, C., et al. (2022). A Large-scale and Longitudinal Measurement Study of DKIM Deployment. Proceedings of the 31st USENIX Security Symposium, 1185 to 1201. www.usenix.org/conference/usenixsecurity22/presentation/wang-chuhan
53. Weichselbaum, L., Spagnuolo, M., Lekies, S., and Janc, A. (2016). CSP Is Dead, Long Live CSP! On the Insecurity of Whitelists and the Future of Content Security Policy. Proceedings of the ACM Conference on Computer and Communications Security (CCS).
54. Whalen, S. checkdmarc: a Python module and command line tool for validating SPF and DMARC DNS records, version 6.0.4. github.com/domainaware/checkdmarc
55. Wilson, E. B. (1927). Probable Inference, the Law of Succession, and Statistical Inference. Journal of the American Statistical Association, 22(158), 209 to 212.
56. Yahoo. Sender Best Practices: Sender Requirements and Recommendations. Enforcement from February 2024. senders.yahooinc.com/best-practices/

Appendices

Appendix A. Headline measures with confidence intervals

MEASURE	K / N	SHARE	95% CI
No enforced DMARC policy	3,790 / 4,402	86.1%	85.0% to 87.1%
Same, among domains that receive mail	2,831 / 3,355	84.4%	83.1% to 85.6%
DMARC reject policy	257 / 4,402	5.8%	5.2% to 6.6%
Pass all five basic controls (fully assessed)	357 / 4,183	8.5%	7.7% to 9.4%
Fail two or more basic controls	2,531 / 4,406	57.4%	56.0% to 58.9%
No http to https redirect	1,059 / 4,186	25.3%	24.0% to 26.6%
No security headers	1,895 / 4,186	45.3%	43.8% to 46.8%
TLS 1.3 negotiated	3,703 / 3,963	93.4%	92.6% to 94.2%
WordPress behind current release	264 / 878	30.1%	27.1% to 33.2%
Lookalike domain with mail records	1,141 / 4,406	25.9%	24.6% to 27.2%
Default selector DKIM key, Google Workspace	188 / 780	24.1%	21.2% to 27.2%
Default selector DKIM key, Microsoft 365	390 / 985	39.6%	36.6% to 42.7%
Unused domain without null SPF and reject	1,034 / 1,047	98.8%	97.9% to 99.3%
Trackers without a privacy link	657 / 1,488	44.1%	41.6% to 46.7%

Appendix B. No enforced DMARC policy, by individual provider

PROVIDER	N	NO ENFORCED DMARC	95% CI
Yahoo	20	100.0%	83.9% to 100.0%
IONOS	56	96.4%	87.9% to 99.0%
GoDaddy	142	94.4%	89.3% to 97.1%
Rackspace	44	93.2%	81.8% to 97.7%
Google Workspace	780	91.8%	89.7% to 93.5%
Zoho	24	91.7%	74.2% to 97.7%
No mail at domain	1,047	91.6%	89.8% to 93.1%
AppRiver (filtering)	45	91.1%	79.3% to 96.5%
Other or self hosted	1,175	90.0%	88.2% to 91.6%
Barracuda (filtering)	27	74.1%	55.3% to 86.8%
Microsoft 365	985	71.5%	68.6% to 74.2%
Mimecast (filtering)	27	48.1%	30.7% to 66.0%
Proofpoint (filtering)	28	25.0%	12.7% to 43.4%

Providers with at least 20 businesses. Small groups carry wide intervals.

Appendix C. Website platforms and certificate issuers

PLATFORM	WEBSITES	CERTIFICATE ISSUER	WEBSITES
Unknown or custom	1,791	Let's Encrypt	2,193
WordPress	1,155	Google Trust Services	863
Wix	332	GoDaddy	565
GoDaddy Website Builder	296	No certificate	261
Squarespace	181	Sectigo Limited	105
Duda	126	DigiCert Inc	63
Shopify	107	ZeroSSL GmbH	42
Weebly	75	Amazon	39
Next.js or React	65	Protocol: TLSv1.3	3,703
Joomla	35	Protocol: TLSv1.2	260
Webflow	12	Protocol: No TLS	223
Drupal	11		

Appendix D. Model 2: HSTS and DMARC enforcement

Logistic regression, n = 3,823 https websites (1,583 sending HSTS). McFadden pseudo $R^2 = 0.222$; $p < 0.001$.

TERM	ODDS RATIO	95% CI	P
Website platform: Custom or unknown	0.084	0.069 to 0.102	< 0.001
Website platform: Self managed	0.047	0.038 to 0.059	< 0.001
Sector: Auto	1.071	0.808 to 1.419	0.633
Sector: Beauty, wellness and fitness	0.631	0.447 to 0.890	0.009
Sector: Food and drink	0.899	0.663 to 1.221	0.497
Sector: Health	0.882	0.680 to 1.143	0.342
Sector: Home services and construction	1.003	0.796 to 1.264	0.980
Sector: Other services	0.947	0.708 to 1.267	0.713
Sector: Retail	0.996	0.729 to 1.360	0.978
DMARC enforced	1.430	1.161 to 1.762	< 0.001

Appendix E. Definitions

No enforced DMARC policy	No DMARC policy of quarantine or reject. The owner gives receivers no instruction to treat failing mail differently; receivers may still filter it.
Reject policy	DMARC p=reject applied to all mail: the strongest DMARC instruction. It does not stop compromised accounts, lookalike domains or display name impersonation.
Live domain	A domain delegated in DNS to name servers other than a known parking or resale service.
DKIM key found	A DKIM public key at the provider's default selector (google, or selector1). A lower bound on key publication; it does not show that mail is signed.
Not assessed	A control that could not be observed: a website that did not answer, or email records no resolver answered for. Neither a pass nor a fail.
Hosted builder	Wix, Squarespace, Shopify, GoDaddy Website Builder, Duda, Weebly or Webflow.
Self managed	WordPress, Joomla or Drupal, which the owner or their contractor must update.
Filtering gateway	Mail routed through Proofpoint, Mimecast, Barracuda or AppRiver.
Unsupported PHP	A PHP release older than 8.2, without security support at the time of measurement.
Vulnerable jQuery	jQuery older than 3.5.0, affected by CVE-2020-11022 and CVE-2020-11023.
Lookalike domain	A domain one plausible typo or glyph substitution away from the business's own, or the same name on another common ending.