ELEVATE
SOLUTIONS

GUIDE · CYBERSECURITY

Why Your Business Needs Email Security

The #1 attack vector

Email is the entry point for most breaches. This guide covers advanced filtering, DMARC, DLP, encryption, and BEC protection for regulated firms — what each layer does, the exact settings that matter, and a deployment checklist you can run this quarter.

01 — THE THREAT

Why email is the #1 vector

Attackers do not break in through your firewall. They log in through your inbox. Email reaches every employee, carries authority by default, and is trusted by the people who approve wires, handle patient records, and hold privileged files.

The numbers are not subtle. CISA states that **"Over 90% of successful cyber-attacks start with a phishing email."** CISA Business email compromise alone produced $2.77B in losses across 21,442 complaints in 2024. FBI IC3 2024 Internet Crime Report Total reported cybercrime losses hit $16.6B in 2024 — a record, up 33% year over year. FBI IC3 2024 Internet Crime Report And the human element was involved in 60% of breaches Verizon DBIR 2025 — which is exactly the element email attacks are built to exploit.

90%+	$2.77B	60%
Successful attacks starting with a phishing email CISA	BEC losses across 21,442 complaints, 2024 FBI IC3 2024	Breaches involving the human element Verizon DBIR 2025

How attacks actually arrive

1. Credential phishing

A login page that looks like Microsoft 365 or Google Workspace, reached from a link in a routine-looking message — a shared document, a voicemail notice, a password-expiry alert. Modern kits run adversary-in-the-middle (AiTM) proxies: the fake page passes your real credentials to the real login service, completes your MFA prompt, and steals the resulting session cookie. The attacker is now signed in as the user, MFA and all. This is why "we have MFA" is necessary but not sufficient.

2. Business email compromise (BEC)

No malware, no link, no attachment — just a plausible message asking someone with payment authority to act. Variants: a lookalike domain one character off from a real vendor (*elevate-solutlons.com*), a spoofed display name on a throwaway mailbox, or a genuinely compromised vendor account sending "updated banking instructions" mid-invoice. Because there is no payload, signature-based filters score these messages as clean. The attack is the sentence, not the file.

3. Malware delivery

With macros blocked by default in Office, delivery shifted to container formats that bypass Mark-of-the-Web: ISO and IMG disk images, password-protected ZIPs (password in the message body), LNK shortcut files, and HTML smuggling — an attachment that assembles the payload in the browser, so nothing malicious crosses the gateway intact. The goal is usually a loader that hands off to ransomware operators or info-stealers days later.

4. Thread hijacking

After compromising one mailbox, the attacker replies inside existing conversations. The message arrives from a real colleague or counterparty, in a thread the victim recognizes, quoting genuine prior content. Reply-chain attacks defeat the single best human heuristic — "do I know this sender?" — and are the standard follow-on move after any account takeover.

Why regulated firms are targeted

Law firms, medical practices, and financial services firms concentrate three things attackers want in one inbox:

WHAT YOU HOLD	WHY IT PAYS
Wire authority	Trust accounts, escrow, settlements, and closings move large sums on email instruction. One diverted wire can exceed a year of ransomware takings from the same firm.
PHI & client data	Health records resell and trigger notification duties. Healthcare breaches average $7.42M — the costliest industry 14 years running. IBM Cost of a Data Breach Report 2025
Privilege & confidentiality	Privileged files and deal information create extortion leverage independent of encryption — attackers can monetize the threat of disclosure alone.

Add public directories: state bar listings, NPI registries, and firm websites publish names, roles, and email formats. An attacker can build a precise target list — managing partner, controller, paralegal handling closings — without touching your network.

02 — THE DEFENSE

What real protection looks like

No single control stops all four attack types above. Effective email security is a stack of layers, each catching what the previous one passes. Here is the layered model, in the order a message encounters it.

LAYER	WHAT IT DOES — AND WHAT TO LOOK FOR
Gateway / API filtering	Scores inbound mail on reputation, content, and authentication results. Two architectures: an MX-record gateway in front of the tenant, or an API-integrated layer that scans inside the mailbox (and can claw back messages post-delivery). API-based tools also see internal-to-internal mail — critical for catching thread hijacking after a takeover.
Link rewriting + detonation	Rewrites every URL so it is re-checked at time of click, not time of delivery. Defeats the standard trick of weaponizing a link hours after the message lands. In Microsoft 365: *security.microsoft.com → Email & collaboration → Policies & rules → Threat policies → Safe Links*. Verify "Apply real-time URL scanning" is on and users cannot click through to the original URL.
Attachment sandboxing	Detonates attachments in a virtual machine before release. In Microsoft 365: *Threat policies → Safe Attachments*; use Dynamic Delivery so the message body arrives while the attachment is scanned. In Google Workspace: *Admin console → Apps → Google Workspace → Gmail → Safety → Attachments*, enable sandboxing of encrypted and anomalous attachment types.
Impersonation detection	Flags lookalike domains, spoofed display names, and first-time senders posing as known contacts. In Defender for Office 365 anti-phishing policies: add your executives and finance staff to *user impersonation protection*, add key vendors to *domain impersonation protection*, and enable *mailbox intelligence* so the system learns real correspondence patterns.
DLP rules	Inspects outbound mail for regulated content — SSNs, account numbers, PHI identifiers, client-matter codes — and blocks, encrypts, or requires justification. Start in audit-only mode for two weeks to tune false positives, then enforce. DLP is also your safety net when a compromised internal account starts exfiltrating by email.
Encryption for regulated content	Policy-driven, not user-optional: transport rules that auto-encrypt when DLP detects PHI or financial data, plus a manual trigger (e.g., subject-line keyword or an "Encrypt" button) for known-sensitive threads. Recipient experience matters — pick a method counterparties can actually open, or they will ask you to resend in plaintext.
MFA + conditional access	The backstop when credentials are phished anyway. Prefer phishing-resistant methods (FIDO2 keys, passkeys) for wire-authority and admin roles — AiTM kits defeat push and SMS, not hardware-bound credentials. Block legacy authentication protocols (IMAP/POP/SMTP basic auth) outright; they bypass MFA entirely. Condition access on managed devices and expected locations.

User reporting loop	A one-click Report Phishing button in every mail client, routed to a queue someone actually triages — with automatic clawback of the same message from every other mailbox that received it. Your users see novel attacks before any vendor signature does; the loop turns that into detection.

WHAT THE BUILT-IN SPAM FILTER MISSES

Default tenant filtering is tuned for bulk spam, not targeted attacks. It routinely passes: BEC messages with no link or attachment (nothing to scan); lookalike domains with clean reputations registered last week; links weaponized after delivery; internal-to-internal mail from a compromised colleague (never crosses the gateway); and vendor-compromise mail that authenticates perfectly — because it really is from the vendor's tenant. Each gap above maps to a layer in the table. That is the point of the stack.

Sequencing advice: authentication (Section 03) and MFA hardening cost the least and close the widest gaps — do them first. Then impersonation and link/attachment policies, then DLP and encryption, then the reporting loop and training. The checklist in Section 04 puts this in order.

03 — AUTHENTICATION

DMARC, SPF, DKIM explained

These three DNS-based standards decide whether the world can tell your real mail from forgeries of your domain. They are cheap, vendor-independent, and most firms have them half-configured — which is worse than it sounds, because a permissive setup actively tells receivers to accept spoofed mail.

SPF — who may send for your domain

A DNS TXT record listing the servers authorized to send mail using your domain in the *envelope sender* (the Return-Path, not the From line users see). Example:

```
v=spf1 include:spf.protection.outlook.com include:_spf.salesforce.com -all
```

The trailing mechanism matters: `-all` means "everything else fails", `~all` means "softfail" (treated as suspicious, not rejected). Hard limit: SPF allows at most **10 DNS lookups** per check. Every `include:` costs at least one, and nested includes count. Exceed 10 and SPF returns *permerror* — your record silently stops working. Firms that have bolted on a CRM, a marketing tool, a ticketing system, and an e-signature platform are usually at or past the limit. Fix by removing dead includes or "flattening" includes into IP ranges (which then need maintenance when vendors change IPs).

DKIM — proof the message wasn't altered

Your mail platform signs each outbound message with a private key; receivers fetch the public key from DNS at `selector._domainkey.yourdomain.com` and verify the signature covers the From, Subject, and body. A valid DKIM signature proves the message originated from a system holding your key and arrived unmodified. Unlike SPF, DKIM **survives forwarding**, because the signature travels inside the message. Enable it on your platform (Microsoft 365: *Defender portal → Email authentication settings → DKIM*, publish the two CNAME records it gives you) and separately in every third-party tool that sends as your domain — each has its own keys and DNS records to publish.

DMARC — the policy that ties them together

SPF and DKIM each validate a technical sender field, but neither checks the **From address users actually see**. DMARC closes that gap with *alignment*: a message passes only if SPF or DKIM passes *and* the validated domain matches the From domain. Your DMARC record (a TXT at `_dmarc.yourdomain.com`) then tells receivers what to do with failures:

```
v=DMARC1; p=none; rua=mailto:dmarc-reports@yourdomain.com; fo=1
```

How the three chain on a single message

CHECK	VALIDATES	DEFEATED BY / BLIND SPOT
SPF	Sending server vs envelope sender	Forwarding (new server isn't in your record); doesn't check the visible From line
DKIM	Signature over content + headers	Mailing lists that modify subject/body; senders you never configured keys for
DMARC	Alignment with the visible From + policy	Nothing it claims to cover — but only as strong as the policy you publish

Note the limit: DMARC stops **exact-domain spoofing**. It does nothing against lookalike domains or display-name tricks — that is what impersonation detection (Section 02) is for. You still want it: it protects your clients and vendors from mail forged in your name, and receivers increasingly junk mail from domains with no DMARC at all.

Rollout: none → quarantine → reject

STAGE	DURATION	WHAT YOU DO
`p=none`	4–6 weeks	Monitor only. Collect aggregate reports, identify every legitimate source failing alignment, fix SPF/DKIM for each.
`p=quarantine`	2–4 weeks	Failures go to junk. Ramp with `pct=25` → `pct=100`. Watch reports for legitimate mail you missed.
`p=reject`	Permanent	Failures are refused outright. Add `sp=reject` for subdomains, and publish `v=spf1 -all` plus `p=reject` on parked domains you own but never send from.

Do not skip straight to reject. The monitoring stage exists because every firm discovers senders it forgot — the billing platform, the scan-to-email copier, the marketing tool a partner signed up for in 2021.

Common failure modes

Forwarders: a recipient's auto-forward breaks SPF (the forwarding server isn't yours). DKIM usually survives, and DMARC needs only one aligned pass — this is why you deploy both, not either. **Mailing lists:** lists that add footers or rewrite subjects break DKIM; most modern lists rewrite the From to compensate. **Third-party senders:** a tool sending "as you" with its own envelope domain passes SPF for *its* domain — unaligned, so DMARC fails unless you set up DKIM with your domain in that tool. **Subdomains:** without `sp=`, your subdomain policy defaults to the organizational policy — verify it covers what you intend.

How to read a DMARC report

Aggregate (rua) reports arrive daily as XML from each major receiver. Raw XML is unreadable at volume — route the rua address into a parsing service and review weekly. Each record gives you: **source IP** (who sent), **count**, **disposition** (what the receiver did), and **SPF/DKIM results with alignment**. Triage rule of thumb: a known vendor IP failing alignment is a configuration task; an unknown IP range sending hundreds of messages as your domain is active spoofing — and proof your move to `p=reject` matters. Forensic (ruf) reports are largely discontinued by major receivers; do not plan around them.

04 — EXECUTION

Deployment checklist

Run top to bottom. Items 1–8 are the authentication track from Section 03; items 9–14 layer on the controls from Section 02. Most firms can complete 1–6 in the first month with no user-visible change.

- [ ] **Inventory every service that sends as your domain** 1
  Mail platform, CRM, marketing, billing, ticketing, e-signature, scanners/copiers, monitoring alerts. Check DMARC reports and accounts-payable records for ones nobody remembers.

- [ ] **Publish a correct SPF record — and stay under 10 lookups** 2
  One TXT record, every legitimate source included, ending in `-all`. Count nested lookups; flatten or prune if over the limit.

- [ ] **Enable DKIM on your primary mail platform** 3
  M365: Defender portal → Email authentication settings → DKIM, publish both CNAMEs, rotate to 2048-bit keys. Google: Admin console → Gmail → Authenticate email.

- [ ] **Enable DKIM in every third-party sender** 4
  Each tool from item 1 gets its own DKIM setup with your domain, so its mail aligns. Tools that can't sign with your domain should send from a subdomain you delegate.

- [ ] **Publish DMARC at p=none with aggregate reporting** 5
  `v=DMARC1; p=none; rua=mailto:…; fo=1` — route rua to a parser, not a human mailbox.

- [ ] **Review reports for 4–6 weeks; fix every legitimate unaligned sender** 6
  Weekly review. Goal: all legitimate volume passing aligned SPF or DKIM before any enforcement.

- [ ] **Tighten to p=quarantine, then p=reject** 7
  Ramp pct 25 → 100 at quarantine; hold 2–4 weeks; then reject. Add `sp=reject` and lock down parked domains with `v=spf1 -all` + reject.

- [ ] **Verify inbound enforcement of other domains' policies** 8
  Confirm your gateway honors DMARC on inbound mail and doesn't override reject dispositions — protection runs both directions.

- [ ] **Configure impersonation and anti-phishing policies** 9
  Protect executives, finance staff, and key vendors by name and domain; enable mailbox intelligence; set first-contact safety tips.

- [ ] **Enable link rewriting and attachment sandboxing** 10
  Time-of-click URL scanning with click-through disabled; attachment detonation with Dynamic Delivery so mail isn't delayed.

- [ ] **Tag external mail and add high-risk transport rules** 11
  External-sender banner; flag or hold messages with banking-change language, lookalike reply-to domains, or executive display names from outside the tenant.

- [ ] **Deploy DLP and policy-based encryption for regulated content** 12

  Rules for SSNs, account numbers, and PHI identifiers; audit-only for two weeks, then enforce with auto-encryption on matched outbound mail.

- [ ] **Enforce MFA, block legacy auth, deploy the report button** 13

  Phishing-resistant MFA for wire-authority and admin roles; conditional access blocking IMAP/POP/basic SMTP; one-click phishing reporting routed to a triaged queue with tenant-wide clawback.

- [ ] **Train, test, and put a quarterly review on the calendar** 14

  Short scenario-based training (BEC and thread hijacking, not just generic phishing); periodic simulations; quarterly re-check of items 1–13, since every new SaaS tool re-breaks item 1.

> **ONE PROCESS CHANGE WORTH MORE THAN ANY FILTER**
>
> Adopt a standing rule: any change to payment instructions — new account, new routing, new payee — is verified by phone at a number you already have on file, never one supplied in the email. Write it into your payment procedures, tell your vendors and clients you follow it, and apply it with no exceptions, including for messages that appear to come from a partner or principal. This single control defeats the costliest email attack outright.

Checklist owner _____

Target completion date _____

Next quarterly review _____

ELEVATE SOLUTIONS

Elevate Solutions provides managed IT and cybersecurity for businesses that answer to regulators, clients, and courts.

Elevate Solutions

6230 Wilshire Blvd Suite 2120, Los Angeles CA 90048

(424) 400-6625 · support@elevatesolutions.io · elevatesolutions.io