CyberFortress

PRODUCT DATASHEET

Mail Gateway

Zero-trust and APT protection for the inbox

Email is still the most reliable way into an organisation – and the attacks arriving through it no longer look like attacks. Mail Gateway analyses traffic end to end, so threats are resolved before they reach your systems.

cyber-fortress.com | info@cyber-fortress.com

THE PLATFORM

Not a standalone gateway

Mail Gateway shares the same core as every other module. A malicious attachment caught in the mail flow immediately becomes an indicator that Threat Hunter, Cloud Hunter and your SIEM all act on – and the endpoint that already received a copy is found in the same motion. Central policy, quarantine and reporting live in MA 360°; behavioural judgement comes from CyberFortress AI. One console, one verdict, one audit trail.

MA 360° — the platform core

Collects telemetry from every security layer, correlates it in real time and orchestrates the response. Mail Gateway hands it every verdict it produces and reads back the shared intelligence that every other module is writing to.

THE FOUR SECURITY MODULES

Mail Gateway

Multi-layered defence against phishing, BEC, ransomware and mail-borne APTs. AI anti-phishing, attachment sandboxing, click-time URL rewrite, inline and BCC deployment, plus an Outlook add-in.

Malware Analyzer 360

Dynamic sandbox detonation, static code inspection, live URL verdicts and Content Disarm & Reconstruction – the engine every attachment and link in the mail flow is handed to.

Threat Hunter

A lightweight endpoint agent backed by central analysis. When a weaponised attachment is found, the endpoints that already received a copy are identified in the same motion.

Cloud Hunter

Container and image analysis with CVE-based scanning, bucket exposure checks and cloud file system protection across hybrid and multi-cloud estates.

Data Layer with CyberFortress TI

Senders, domains and links are checked against the same live intelligence the rest of the platform uses – terabytes a day, correlated in real time.

CyberFortress AI

Detects impersonation, BEC and targeted spear phishing by behaviour rather than content, and learns what normal correspondence looks like in your organisation.

WHY MAIL GATEWAY

The door attackers still knock on

Most email security starts from the wrong assumption: that a message is fine unless something known-bad is found in it. Zero trust inverts that. Nothing is delivered on the strength of looking legitimate – every sender, every attachment and every link has to earn it.

That distinction matters because the attacks that succeed are the ones built to look legitimate. A Business Email Compromise message carries no attachment and no link. A spear-phishing campaign is written for one company, sometimes one person, and has never been seen anywhere else. A weaponised document is a real document until the macro runs.

Mail Gateway verifies continuously rather than once at the perimeter. Attachments are detonated in the sandbox, links are re-evaluated at the moment they are clicked, and sender behaviour is weighed against what normal correspondence in your organisation actually looks like. That is what makes it effective against zero-day attacks carrying no signature on the day they arrive.

FOUR ATTACKS, ONE CHANNEL

Phishing
Credential-harvesting pages behind convincing branding and shortened links.

Business Email Compromise
No attachment, no link – just a convincing request for a payment or a password.

Ransomware
Weaponised documents and droppers, detonated before delivery.

Spear phishing & APT
Campaigns built for one company, caught by behaviour rather than signature.

99.9%
AI anti-phishing detection accuracy on fraudulent messages

Zero
Messages trusted by default – every one is verified

On click
Links re-verified at click time, not just on delivery

99.99%
Targeted availability, so mail flow is never the bottleneck

What survives each layer

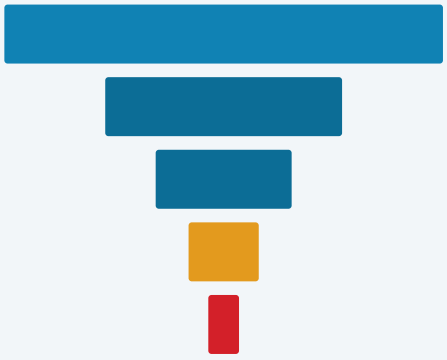

Inbound messages
100% of traffic

Past reputation and policy checks
~54%

Past spam and bulk filtering
~31%

Attachments and links to analysis
~9%

Blocked or rebuilt with CDR
~1.4%

Indicative shape of a corporate mail flow. Bulk filtering clears the noise that targeted attacks rely on to blend in; only what is left is worth the cost of detonation.

INSIDE THE GATEWAY

Six layers, one mail flow

Four attack types share one delivery channel – and are handled by different layers of the same engine. Every layer writes what it learns back to the core.

AI anti-phishing

Sender reputation, header anomalies, language patterns, brand impersonation and visual similarity are weighed together rather than checked one by one. The model builds a picture of what normal correspondence looks like in your organisation, so a well-crafted fake stands out even when it breaks no fixed rule. **99.9% detection accuracy.**

Attachment sandboxing

Every attachment is handed to the Malware Analyzer 360 and executed in isolation, so both known families and zero-day samples are judged on behaviour. Where the business cannot wait for a verdict, Content Disarm & Reconstruction delivers a rebuilt, safe-to-open version of the document immediately.

URL rewrite

Links in delivered mail are rewritten so the destination is re-evaluated at the moment the user clicks – not only at the moment the mail arrived. A page that was harmless on delivery and turns malicious hours later is quarantined before the interaction completes, which is the standard trick in targeted phishing.

Inline & BCC analysis

Inline mode inspects mail in transit and blocks threats on the spot. BCC mode observes a copy of the traffic, giving full visibility with zero risk to delivery. The architecture is built to detect and block without introducing the latency that usually comes with in-path scanning.

Spam & bulk filtering

High-volume unwanted mail is separated out before it reaches users, keeping the signal-to-noise ratio high for both people and analysts. Clearing the bulk traffic also removes the cover that targeted attacks rely on to blend into a crowded inbox.

Outlook add-on

Directly inside Microsoft Outlook, users can submit an attachment for analysis, check a link's real destination, or start a spam review on a suspicious message. Turning "this looks odd" into a real submission takes one click, so your users become a working detection layer instead of a reporting backlog.

Two ways to deploy, one engine

INLINE MODE – inspects in transit and blocks on the spot.

01 **Inbound**
Mail arrives at the gateway in path, before any mailbox sees it.

02 **Analyse**
Sender, body, links and attachments are resolved together.

03 **Deliver**
Delivered clean, rebuilt with CDR, or held in quarantine.

BCC MODE – full visibility with zero risk to delivery.

01 **Inbound**
Mail is delivered exactly as it is today.

02 **Copy**
A BCC copy is streamed to the analysis engine.

03 **Report**
Findings only – the mail flow is never touched.

Most estates start in BCC mode alongside an existing gateway, then move in path once the findings are understood. Both modes run the same analysis and report into the same console.

ZERO TRUST IN PRACTICE

Verified again at the moment it matters

Delayed weaponisation is the reason delivery-time scanning stops working. The attacker sends a link to a page that is genuinely harmless, waits for it to clear the gateway, and swaps the content hours later – after every check has already passed.

A delayed-weaponisation campaign, hour by hour

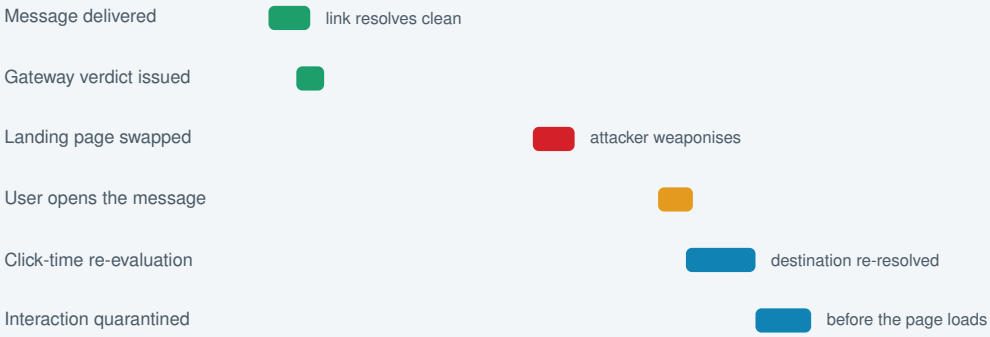

Delivery-time scanning sees only the first hour. Click-time re-evaluation is what covers the rest of the window.

What a message has to earn

Sender

Reputation, header anomalies, authentication records and behavioural history are weighed together – impersonation and unusual sender patterns are flagged before content is even read.

Attachment

Detonated in an isolated sandbox and read at code level. Where delivery cannot wait, CDR strips active content and rebuilds the document clean.

Link

Resolved through the full redirect chain at delivery, then resolved again at click time – the destination is judged on where it actually goes.

What happens to the verdict

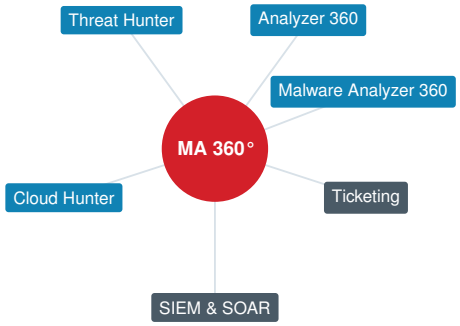

Message verdict	deliver, rebuild or hold
Extracted IOCs	senders, domains, hashes
Estate sweep	who else received it
Audit trail	one record, one console

IN PRACTICE

Where teams put it to work

The gateway is usually the first module switched on, because the evidence it produces is immediate and the deployment carries no risk to mail flow.

Evaluating what gets through today

BCC mode alongside the existing gateway for 30 days produces a report of exactly what the current filter missed – with the samples, the senders and the verdicts attached.

Replacing delivery-time-only scanning

Click-time URL re-evaluation closes the window that delayed weaponisation depends on, which is the gap most legacy gateways still leave open.

Turning user reports into verdicts

Phish Control takes a reported message from there: the attachment is detonated, the links are resolved and a spam review is started, without an analyst touching it.

Executive and finance protection

Behavioural analysis flags impersonation and unusual sender patterns that content scanning cannot see – the BEC case, where there is nothing to scan.

Regulated correspondence

Full audit trail of every verdict, every release and every administrative action, retained in a form that satisfies a regulator asking what was delivered and why.

Estate-wide containment

When an attachment is judged malicious after delivery, the endpoints that already received a copy are identified and swept in the same motion.

Detection accuracy by attack type

Indicative figures across an internal corpus. BEC carries no payload to scan, which is why behavioural analysis rather than content matching drives the result.

Find out what your current filter is missing.

Deploy in BCC mode alongside your existing gateway for 30 days. No change to mail flow, no risk to delivery – just a report of what got through.

cyber-fortress.com
info@cyber-fortress.com

SPECIFICATION

Features and capabilities

THREAT COVERAGE

- Phishing and credential-harvesting campaigns
- Business Email Compromise and payment fraud
- Ransomware droppers and weaponised documents
- Spear phishing and targeted APT campaigns
- Mail-borne zero-day and previously unseen samples
- Brand and executive impersonation
- Display-name and look-alike domain spoofing
- Malicious QR codes and image-embedded payloads
- Thread hijacking and reply-chain abuse

AI ANTI-PHISHING

- 99.9% detection accuracy on fraudulent messages
- Sender reputation and behavioural history scoring
- Header anomaly and authentication record analysis
- Language and tone pattern analysis
- Brand impersonation and visual similarity detection
- Learns the normal correspondence pattern of your organisation
- Proactive detection of targeted fraud and identity spoofing
- Classification in milliseconds, not minutes

ATTACHMENT ANALYSIS

- Real-time detonation in the Malware Analyzer 360 sandbox
- Static code inspection alongside dynamic behaviour
- All file types, including nested and password-protected archives
- Office documents, PDFs, scripts, installers and executables
- Content Disarm & Reconstruction rebuilds a clean, usable file
- Macros, embedded scripts and malformed structures stripped
- Clean copy delivered immediately where the business cannot wait
- Verdicts carry extracted IOCs and MITRE ATT&CK mapping

URL PROTECTION

- Links rewritten for click-time re-evaluation
- Full redirect chain resolved to the real destination
- Shorteners and redirectors judged on where they end
- Delayed weaponisation caught after delivery
- Phishing pages, C2 endpoints and drive-by chains identified
- Interaction quarantined before the page loads
- Live URL verdicts shared with the rest of the platform

SPAM & BULK FILTERING

- High-volume unwanted mail separated before it reaches users
- Newsletter and bulk traffic classified separately from spam
- Per-user and per-domain allow and block policy
- User-initiated spam review from inside Outlook
- Clears the cover that targeted attacks blend into

DEPLOYMENT & MAIL FLOW

- Inline mode – inspects in transit and blocks on the spot
- BCC mode – observes a copy with zero risk to delivery
- Built to block without in-path scanning latency
- 99.99% targeted availability
- On-premise, cloud or hybrid deployment
- SMTP, IMAP and Postfix mail flow integration
- Runs alongside an existing gateway during evaluation
- Sits in front of Microsoft 365, Google Workspace or on-prem mail

USER TOOLS

- Dedicated Microsoft Outlook add-in
- Submit an attachment for analysis in one click
- Check a link's real destination before opening it
- Start a spam review on a suspicious message
- Reporting users become a working detection layer
- Plain-language explanation of why a message was held

POLICY & QUARANTINE

- Central policy, quarantine and reporting in MA 360°
- Per-user, per-group and per-domain rule sets
- Release, delete or re-analyse held messages
- Retrospective removal of already-delivered messages
- Role-based access to quarantine and analysis results
- Full audit trail of every verdict and administrative action

REPORTING & NOTIFICATIONS

- Customisable logs streamed to your syslog server
- Analysis results emailed in any specified format
- Dynamic, user-friendly reports distributed to third parties
- PDF, CSV, JSON and HTML report generation
- Health monitor showing the status of each engine
- Real-time disk capacity monitoring

INTEGRATIONS

- RESTful API with bi-directional data and action flow
- SIEM: Splunk, Sentinel, QRadar, Elastic, LogRhythm, ArcSight
- SOAR: Cortex XSOAR, Splunk SOAR, Forti-SOAR, Tines, Swimlane, Torq
- EDR/EPP: CrowdStrike Falcon and Microsoft Defender
- Threat intelligence: MISP, OpenCTI, ThreatQ, Anomali ThreatStream
- Ticketing systems open, update and close incidents automatically
- Syslog streaming for everything else

CyberFortress

No message is trusted by default.

Mail Gateway is one of four modules on a single core. Deploy it on its own, or switch on the layers around it and let one verdict drive the whole estate.

THE PLATFORM AT A GLANCE

MA 360°
The core — collection, correlation and orchestration.

Malware Analyzer 360
Sandbox, static analysis, URL verdicts and CDR.

CyberFortress TI
Comprehensive Threat intelligence and analytics.

Mail Gateway
Phishing, BEC, ransomware and mail-borne APTs.

Threat Hunter
Lightweight agent, YARA sweeps and air-gap modes.

Cloud Hunter
Container, bucket and cloud file system protection.

CyberFortress AI
Behavioural analytics wrapping every module.

cyber-fortress.com

info@cyber-fortress.com