The following documentation is an electronically-submitted vendor response to an advertised solicitation from the *West Virginia Purchasing Bulletin* within the Vendor Self-Service portal at ***wvOASIS.gov***. As part of the State of West Virginia's procurement process, and to maintain the transparency of the bid-opening process, this documentation submitted online is publicly posted by the West Virginia Purchasing Division at ***WVPurchasing.gov*** with any other vendor responses to this solicitation submitted to the Purchasing Division in hard copy format.

Header 1

List View

General Information | Contact | Default Values | Discount | Document Information | Clarification Request

Procurement Folder: 2028979
Procurement Type: Central Master Agreement
Vendor ID: VC0000027508
Legal Name: CAMPUSGUARD LLC
Alias/DBA:
Total Bid: $1,514,040.00
Response Date: 08/24/2026
Response Time: 16:28
Responded By User ID: jhaack
First Name: Jennifer
Last Name: Haack
Email: jhaack@campusguard.com
Phone: 6122076336

SO Doc Code: CRFQ
SO Dept: 0231
SO Doc ID: OOT2700000005
Published Date: 8/18/26
Close Date: 8/25/26
Close Time: 13:30
Status: Closed
Solicitation Description: Addendum No 1-Advanced Email Security
Total of Header Attachments: 1
Total of All Attachments: 1

Department of Administration
Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

State of West Virginia
Solicitation Response

Proc Folder:	2028979
Solicitation Description:	Addendum No 1-Advanced Email Security
Proc Type:	Central Master Agreement

Solicitation Closes	Solicitation Response	Version
2026-08-25 13:30	SR 0231 ESR08242600000001284	1

VENDOR

VC0000027508
CAMPUSGUARD LLC

Solicitation Number: CRFQ 0231 OOT2700000005

Total Bid: 1514040 **Response Date:** 2026-08-24 **Response Time:** 16:28:10

Comments:

FOR INFORMATION CONTACT THE BUYER
Toby L Welch
(304) 558-8802
toby.l.welch@wv.gov

Vendor
Signature X **FEIN#** **DATE**

All offers subject to all terms and conditions contained in this solicitation

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
1	Core Detection	22000.000	EA	16.450000	361900.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments:

Extended Description:

Core Detection per Mailbox estimated 22,000 Mailboxes
The usage amount provided is an estimate only. Actual usage may be higher or lower, and the figure is not guaranteed.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
2	Core Detection - Year 2 (Optional Renewal)	22000.000	EA	16.940000	372680.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments:

Extended Description:

Core Detection per Mailbox estimated 22,000 Mailboxes
The usage amount provided is an estimate only. Actual usage may be higher or lower, and the figure is not guaranteed.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
3	Core Detection - Year 3 (Optional Renewal)	22000.000	EA	17.450000	383900.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments:

Extended Description:

Core Detection per Mailbox estimated 22,000 Mailboxes
The usage amount provided is an estimate only. Actual usage may be higher or lower, and the figure is not guaranteed.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
4	Core Detection - Year 4 (Optional Renewal)	22000.000	EA	17.980000	395560.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments:

Extended Description:

Core Detection per Mailbox estimated 22,000 Mailboxes
The usage amount provided is an estimate only. Actual usage may be higher or lower, and the figure is not guaranteed.

CAMPUSGUARD®

Response to Request for Quotation

CRFQ 0231 OOT2700000005
Advanced Email Security

Prepared for:

Toby L. Welch
West Virginia Office of Technology
304.558.8802
Toby.L.Welch@wv.gov

Prepared by:

Simon Denis
Director of Business Development | Midwest
CampusGuard
231.408.8009
sdenis@campusguard.com

Table of Contents

Department of Administration
Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

State of West Virginia
Centralized Request for Quote
Info Technology

Proc Folder: 2028979 **Doc Description:** Advanced Email Security **Proc Type:** Central Master Agreement			**Reason for Modification:**
Date Issued	**Solicitation Closes**	**Solicitation No**	**Version**
2026-07-31	2026-08-19 13:30	CRFQ 0231 OOT2700000005	1

BID RECEIVING LOCATION

BID CLERK
DEPARTMENT OF ADMINISTRATION
PURCHASING DIVISION
2019 WASHINGTON ST E
CHARLESTON WV 25305
US

VENDOR

Vendor Customer Code:
Vendor Name : CampusGuard LLC
Address : 121 South 13th St, Suite 201
Street :
City : Lincoln
State : Nebraska **Country :** USA **Zip :** 68508
Principal Contact : Simon Denis
Vendor Contact Phone: 231.408.8009 **Extension:**

FOR INFORMATION CONTACT THE BUYER
Toby L Welch
(304) 558-8802
toby.l.welch@wv.gov

Vendor Signature X **FEIN#** 264478556 **DATE** 8/11/26

All offers subject to all terms and conditions contained in this solicitation

ADDENDUM ACKNOWLEDGEMENT FORM
SOLICITATION NO.: CRFQ OOT27*005

Instructions: Please acknowledge receipt of all addenda issued with this solicitation by completing this addendum acknowledgment form. Check the box next to each addendum received and sign below. Failure to acknowledge addenda may result in bid disqualification.

Acknowledgment: I hereby acknowledge receipt of the following addenda and have made the necessary revisions to my proposal, plans and/or specification, etc.

Addendum Numbers Received:
(Check the box next to each addendum received)

[X]	Addendum No. 1	[ ]	Addendum No. 6
[ ]	Addendum No. 2	[ ]	Addendum No. 7
[ ]	Addendum No. 3	[ ]	Addendum No. 8
[ ]	Addendum No. 4	[ ]	Addendum No. 9
[ ]	Addendum No. 5	[ ]	Addendum No. 10

I understand that failure to confirm the receipt of addenda may be cause for rejection of this bid. I further understand that any verbal representation made or assumed to be made during any oral discussion held between Vendor's representatives and any state personnel is not binding. Only the information issued in writing and added to the specifications by an official addendum is binding.

CampusGuard LLC
Company

Authorized Signature

8/24/2026
Date

NOTE: This addendum acknowledgement should be submitted with the bid to expedite document processing.

Executive Summary

The West Virginia Office of Technology seeks advanced email security layered on its existing Google Workspace environment. The proposed solution delivers behavioral detection of phishing, business email compromise, and account takeover through the Google Workspace API, without an MX-record change, mail-flow redirection, or intermediary gateway.

This architecture is a direct response to Section 4.1.1.1.1. Mail is delivered natively by Google Workspace and analyzed immediately afterward through authorized API access. Because analysis is not in the delivery path, there is no new point of failure in State mail flow and no latency introduced into message delivery. Detection is behavioral rather than rule-based, and applies to internal State-to-State mail on the same terms as external mail — the surface a gateway positioned at the perimeter does not observe.

Where Agency policy directs it, a malicious message is quarantined within seconds of delivery across all affected mailboxes, satisfying the automated remediation required by Section 4.1.5. Real-time, time-of-click URL inspection (Section 4.1.3.2) is provided as a server-side capability the State can enable, without any change to MX records or mail flow.

Attachments are subjected to deep file inspection and to cloud-based sandboxing with dynamic detonation, as required by Section 4.1.3.3 and as clarified by the State in Addendum No. 1. Analysis is staged so that known and structurally identifiable threats are resolved in seconds, and dynamic detonation is reserved for content that remains genuinely unknown.

Retrospective scanning of existing mailbox content begins immediately upon activation, establishing both the threat picture and the behavioral baseline the detection engine evaluates against, as described in Section 4.1.2.3.

The State contracts with a single vendor for the entirety of Section 4.1. Detection, attachment analysis, data-loss prevention, remediation, administration, and reporting are delivered as one integrated platform, under one contract, with one accountable party. No capability in this response is contingent on a separate award or a second vendor relationship with the State.

The platform is hosted within Nelnet's FedRAMP-authorized AWS environment, satisfying Verification Pathway A under Section 4.1.4, with all customer data stored and processed within the United States.

Pricing is a flat annual subscription of $16.45 per mailbox, inclusive of all updates, tier-3 support, and incident response automation. A firm per-mailbox rate is quoted separately for the base year and for each of the three Optional Annual Renewal Years on the Pricing Pages.

4.1.2 Workspace Integration and Historical Scanning

The solution operates exclusively through Google Workspace APIs, inside the perimeter as defined in Section 3.19, without altering mail routing or MX records. Internal-to-internal traffic across all domains and organizational units is analyzed as described in Section 4.1.2.1, and native Google Workspace and Gemini functionality is preserved as described under Coexistence with Google Workspace and Gemini.

Historical scanning begins immediately upon activation rather than on a scheduled window. Enrollment and retrospective processing are concurrent: as each mailbox completes enrollment it enters the retrospective queue, so historical analysis is underway across the earliest mailboxes while later mailboxes are still enrolling. No State action initiates it beyond the single domain-wide delegation authorization described in Section 4.1.6.3.

4.1.2.3 Historical Scanning

Retrospective scanning begins immediately upon activation, as required by Section 4.1.2.3, and runs to completion without any change to mail routing and without interruption to active State communications.

Scope of the pass. Scanning covers all existing mailbox content across every domain and organizational unit in the tenant — messages, attachments, and the current active state of mailbox rules — for every mailbox enrolled during onboarding, including shared, delegated, resource, and service-account mailboxes. Consistent with the State's clarification, mailbox rules are read as their present configuration — forwarding, deletion, auto-move, auto-archive, delegation, and auto-reply settings in effect on each mailbox — rather than as a forensic history of rule changes. Consistent with the State's clarification that initial retroactive scanning upon implementation may be scoped to approximately one year or less, the initial window is agreed with the State at implementation; the platform supports a window of a year or more where the State elects it, and the trailing ninety days are processed first so the behavioral baseline is complete and evaluable at the earliest point in the engagement.

The pass produces two distinct results. Retrospective scanning is not a single sweep with a single output. One enumeration of historical content feeds two independent consumers, and both are required.

First — threats already resident in the environment. Malicious content delivered before monitoring began is still in the mailbox and still dangerous. Historical messages receive the same impersonation, look-alike, reply-to, authentication-result, and URL-reputation analysis as newly delivered mail. Historical attachments are analyzed by file-hash threat intelligence, signature scanning, and static structural inspection — macros, embedded scripts, PDF actions, archive and container contents, and filename evasion. Any finding is raised and dispositioned under Agency policy on the same terms as a live detection, and carries

provenance identifying it as a historical finding so an administrator can distinguish content discovered in the archive at onboarding from mail arriving today. Because a first pass over a mature environment characteristically surfaces a substantial set of findings at once, results are presented as a reviewable set rather than as an undifferentiated alert stream.

Second — the behavioral baseline the detection engine runs on. This is the durable output, and it is the reason historical processing matters beyond the one-time sweep. Detection is behavioral: a message is judged by how far it departs from the State's established correspondence patterns. Those patterns are learned from history. Processing the historical window populates, for every correspondent relationship in the environment, the direction and volume of correspondence, its cadence and the typical interval between messages, the hours of day and days of week on which it normally occurs, and typical message size. Impersonation, look-alike, dormancy, first-contact, and send-time-outlier analysis all resolve against those learned baselines.

Why this changes day-one detection quality. Without a learned history, a behavioral engine must treat every correspondent as unfamiliar. The practical consequence is not merely weaker detection — it is a specific and predictable class of false positive. A financial request from a long-established State vendor is indistinguishable, to a system with no history, from a financial request from a stranger; a routine monthly invoice from a correspondent the State has transacted with for years reads as first contact. Processing history first is what prevents that. Equally, it is what makes genuine anomalies legible: a correspondent that has been dormant for eight months and suddenly resumes contact with a payment-detail change is only identifiable as dormant if the eight months of silence were observed.

The baseline persists and continues to improve. Historical processing is not a transient warm-up discarded after onboarding. The relationship model it establishes is durable and is extended continuously by live mail thereafter, so the State's baseline deepens over the life of the contract rather than resetting. Analysis performed at any later point in the engagement is evaluated against the full observed history — the retrospective window plus everything since — not against a rolling recent sample.

Operational properties. Retrospective scanning executes on a separate processing path from live monitoring and is rate-governed against Google Workspace API quota, so backfill never delays analysis of newly delivered mail and never consumes tenant API headroom in a way that would affect the State's other Workspace integrations. Live monitoring always takes precedence for a contended request slot. Progress is checkpointed per mailbox, so an interruption resumes rather than restarts. Historical processing reads mailbox content but retains only derived relationship statistics and the evidence supporting a finding; message bodies and attachment content are not retained beyond analysis except where a finding's disposition requires preserving the original, which follows the standard encrypted-quarantine path.

Completion reporting. On completion a scan report is delivered to the State recording mailboxes processed and reconciled against the Workspace directory, the historical window actually covered per mailbox, message and attachment volume analyzed, findings by category, the proportion of correspondent relationships for which a complete baseline was established, and any mailbox that could not be fully processed together with the reason.

Historical scanning introduces no mail-routing change, no MX change, and no maintenance window. Mail delivery continues uninterrupted throughout.

4.1.3 Technical & Threat Protection Capabilities

Detection combines machine learning with behavioral baselining and anomaly detection, rather than static rule-based filtering, as required by the Section 4.1.3 preamble. Language-pattern analysis is performed by a machine-learning classifier that evaluates message intent — identifying payment, payroll-redirection, and credential-solicitation intent independently of whether the wording matches any known campaign. Behavioral baselining and anomaly detection operate on the learned per-relationship model described in Section 4.1.2.3: each message is scored against the correspondence patterns actually observed in the State's environment, and multiple weighted signals are evaluated together to produce a single judgment rather than applied as independent pass/fail rules. The historical corpus processed at activation is what that model is learned from, so behavioral detection is operative from the beginning of the engagement rather than after a warm-up period.

Analysis covers every message delivered to a monitored mailbox — mail originating outside the State and mail exchanged between State mailboxes alike — across all domains, agencies, and organizational units. Internal-to-internal traffic is evaluated by the same engine, and held to the same evidentiary standard, as externally originated mail, as required by Section 4.1.2.1 and the definition at Section 3.16.

High-confidence findings are acted upon in accordance with Agency policy. When policy elects quarantine, a threat is typically quarantined within seconds of delivery.

Position Relative to Google's Native Security Controls

The Agency's stated objective is the detection and prevention of threats that bypass native security controls. The proposed solution is a complementary layer, not a replacement. Google Workspace continues to apply its native controls at delivery, and the platform consumes their output rather than duplicating it — Google's SPF, DKIM, and DMARC authentication results are read as inputs to detection, not re-derived. Analysis is directed at the categories that characteristically survive native filtering:

- **Threats carrying no malicious artifact.** A payment-redirection or payroll-diversion request from a look-alike domain contains no attachment, no malicious link, and no signature-detectable content. Nothing about it is technically malicious; what identifies it is that it departs from an established relationship. Reputation and signature filtering has no purchase on it.
- **Internal-to-internal mail.** Mail exchanged between State mailboxes is delivered inside the tenant and never crosses the perimeter, so a control positioned at the perimeter never evaluates it. Once an attacker holds valid State credentials, this is the surface they use.

- **Compromised legitimate senders.** A message from a genuine, authenticated correspondent whose account has been taken over passes every authentication check, because the authentication is real. It is identifiable only against that correspondent's established behavior.
- **Threats weaponized after delivery.** A link benign at delivery and malicious hours later was correctly cleared by delivery-time inspection. Continuous post-delivery re-evaluation is the only control that catches it.

Each of these is addressed by behavioral analysis against a learned baseline rather than by a second pass of the same signature and reputation checks Google has already applied.

4.1.3.1 Business Email Compromise (BEC) & Spoofing

BEC detection covers internal and external impersonation, vendor compromise, invoice fraud, and look-alike domain analysis, as required by Section 4.1.3 and Section 4.1.3.1. The following signals are evaluated together and scored, rather than applied as independent rules:

- **Communication baselining.** Learned per-sender behaviour — established correspondents, sending cadence, and typical send times. Departure from that baseline is a primary signal, including a long-dormant account that begins sending and messages sent well outside a sender's established pattern. Communication baselines are established from the historical corpus processed at activation (Section 4.1.2.3) and extended continuously by live mail thereafter, so relationship-dependent signals are operative from the start of the engagement rather than accumulating from the deployment date forward.
- **Language-pattern and intent analysis.** Classification of message intent — in particular financial-request intent — identifying the payment, payroll-redirection, and credential-solicitation requests characteristic of business email compromise, independently of whether the wording matches a known campaign.
- **Executive and display-name impersonation.** Display-name spoofing evaluated against the organizational directory, including internal impersonation from a compromised or external account.
- **Look-alike and homoglyph domain analysis.** Registrable-domain comparison against known-good correspondents, covering character substitution, homoglyphs, insertion, and transposition.
- **Reply-To redirection.** Divergence between the visible sender and the address that would receive a reply.
- **Vendor compromise and invoice fraud.** Payment-detail and banking-instruction change requests, and financial solicitations arriving from a correspondent with no established history with the State.
- **Authentication breaks.** Messages whose authentication result is inconsistent with the identity they assert, including spoofing of senders with which the State has an established relationship. SPF, DKIM, and DMARC results are read from the

authentication results Google records at delivery, so authentication-dependent signals evaluate correctly without the platform participating in the SMTP session.

These signals apply to internal State-to-State mail on the same terms as externally originated mail. That matters specifically for business email compromise: once an attacker holds valid State credentials, the messages they send to colleagues originate inside the tenant and carry a legitimate sender identity, so impersonation and financial-request signals must be evaluated on internal traffic to catch them at all.

A BEC or account-takeover message is detected and, if Agency policy says so, quarantined within seconds of delivery. The proposed solution does not sit in the delivery path and does not claim to block a message before Google delivers it.

4.1.3.2 Advanced Phishing & Credential Harvesting

URLs are captured from the message body and headers after delivery and scored against multiple reputation sources, with redirect chains followed to their final destination so that a benign opening hop does not conceal a malicious endpoint. Because phishing links are frequently benign when delivered and weaponized hours later, every URL in delivered mail is continuously re-evaluated against updated threat intelligence after delivery; when a URL is later determined malicious, the message containing it is automatically removed from every mailbox that received it under Section 4.1.5.1, without administrator intervention. This provides continuous protection against links weaponized after the message arrives.

URL extraction covers links in the message body, in message headers, and in attachments, and includes URLs encoded in QR images, which are decoded and scored on the same terms as a text link — a delivery method chosen specifically to defeat text-based link inspection. Redirect chains are followed to their true final destination, and the destination itself is analyzed for credential-harvesting behavior, so a phishing page generated dynamically at request time is judged by what it presents rather than by whether its address is already known.

In addition, the proposed solution provides real-time, time-of-click URL inspection as defined in Section 3.13 — a server-side capability the State can enable. When enabled, at the moment a user activates a link, its destination is evaluated in real time — multi-source reputation, full redirect-chain resolution to the true final destination, and analysis of the destination for credential-harvesting and phishing behavior — and the user is permitted, cautioned with an interstitial, or blocked, according to Agency policy. This delivers inspection at the moment of the click, and not only at the time of delivery.

This capability is delivered entirely from the server side, through the Google Workspace API, after delivery. It requires no email client, no browser extension, and no endpoint software of any kind on any State-managed device, satisfying Section 4.1.1.1.3. It does not route mail through a gateway, relay, or proxy, and it does not change the State's MX records or mail

flow — so the architecture verified under Section 4.1.6.1 is unchanged and Section 4.1.1.1.1 continues to hold. Google remains the mail system of record; delivery to the recipient mailbox is complete before any processing begins.

The proposed solution preserves the delivered message and its transparency:

- **The original message is preserved first.** Before any change, the delivered message is captured intact and encrypted, so the original is always recoverable and can be released to the mailbox unchanged at any time.
- **The original destination is not removed.** The proposed solution does not strip the link's destination; it wraps the link so that activation passes through the inspection step, with the original destination preserved within the link and visible to the user. A user who hovers over a link sees its true destination, and any reviewer can read the original target directly from the message.
- **Authentication is preserved and documented.** The sender's SPF, DKIM, and DMARC results are verified before any modification and recorded with the message; the message placed in the mailbox is re-signed and carries an ARC seal documenting that pre-modification authentication state, so the message presents a valid signature and a complete, auditable authentication history. The original message — including the sender's original signature — is retained intact and independently verifiable.
- **The inspection layer fails open and is never a single point of failure.** Mail continues to be delivered by Google, and delivered mail remains readable, regardless of the inspection service's state. The inspection service is deployed redundantly. If a verdict cannot be produced at the moment of a click, the user is forwarded to the intended destination rather than blocked. Because the intended destination is preserved within the link, the inspection layer never separates the State from a legitimate destination.
- **Only the link target is affected.** Message content, structure, and rendering are otherwise unchanged, so native Google Workspace features — including Gemini smart replies, AI-generated summaries, and native message rendering — continue to operate on the message, as required by Sections 4.1.2.2 and 3.17.

The only change this capability introduces to a delivered message is the link target. It is the minimal change that allows inspection at the moment of the click — required by Section 4.1.3.2 and defined in Section 3.13 — without installing software on State devices. It introduces no mail-flow gateway and no MX change, and the architecture verified under Section 4.1.6.1 remains unchanged.

4.1.3.3 Malware & Zero-Day Attachments

Attachments are retrieved through the Google Workspace API after delivery and subjected to **deep file inspection** and **cloud-based sandboxing with dynamic detonation**, as required by Section 4.1.3 and Section 4.1.3.3. Coverage spans executables, archives, scripts, Office documents, PDFs, and containerized payloads.

Analysis is **staged**. Each stage is capable of returning a conclusive verdict on its own, and an attachment advances to the next stage only when the preceding ones have not resolved it. Dynamic detonation is the terminal stage, reached by what remains genuinely unknown.

- **Stage 1 — Known-file intelligence.** File-hash threat intelligence and a cross-tenant verdict cache. Mail is highly repetitive — signatures, newsletters, recurring reports, the same vendor statement to forty recipients — so a substantial share of attachments have an existing verdict and are resolved immediately.
- **Stage 2 — Signature scanning.** Known-malware signature detection (ClamAV, in-boundary) across all retrieved attachments.
- **Stage 3 — Static structural inspection.** The file's structure is parsed without executing it: Office document macros, embedded scripts and auto-execution triggers; PDF embedded JavaScript, launch actions, embedded files and obfuscated object streams; nested archive expansion, container structure inspection and detection of payloads smuggled inside otherwise valid container formats; and evasion detection covering right-to-left override, double extensions, content-type mismatch, and file types disguised as benign formats.
- **Stage 4 — Relationship and provenance context.** An attachment arriving from a correspondent with an established sending history to the State, of a type that correspondent routinely sends, carries different prior risk than the same file type from a first-contact sender. This context informs what advances rather than deciding the verdict alone.
- **Stage 5 — Dynamic detonation.** What the preceding stages cannot resolve is detonated in an instrumented cloud sandbox and observed while executing: process creation, file and registry activity, code injection and persistence, network callout and command-and-control behavior, and the anti-analysis behavior characteristic of a payload that detects it is being observed. This identifies zero-day and previously unseen payloads by what they do, with no prior signature, hash, or reputation.

Why the Analysis Is Staged

Detonation is the most capable form of attachment analysis and also the most expensive one. A sandbox executes a file in real time, so a verdict takes minutes rather than milliseconds, and sandbox capacity is finite and serialized in a way that signature and structural analysis are not. An architecture that detonates every attachment indiscriminately does not gain protection — it loses it. Capacity saturates, the queue lengthens, and verdict latency stretches for every file including the one that actually mattered.

Staging produces two concrete benefits for the State:

- **Most attachments are resolved in seconds, not minutes.** A file convicted by signature, cleared by hash intelligence, or convicted by structure does not need to

be executed to reach a verdict, and waiting for a sandbox to confirm what is already known would only delay the response.

- **Sandbox capacity is reserved for genuinely unknown content.** Because the upstream stages absorb the volume, detonation capacity is available when an actual unknown arrives — including during a malspam wave, which is precisely when an indiscriminate architecture is most degraded.

This is the same principle applied by mature advanced-threat platforms, where detonation sits at the end of a funnel whose earlier layers exist to keep work off the sandbox. The Vendor's position is not that the earlier stages substitute for detonation — Section 4.1.3.3 requires detonation, and the platform provides it. The position is that detonation performs best when it is reached deliberately.

Verdicts escalate only. A dynamic verdict can raise an attachment's disposition and never lowers one already reached by signature or structure. Attachments that cannot be opened for analysis — password-protected archives, unsupported or oversized files — are treated as unresolved rather than clean, and dispositioned under Agency policy.

No mail-flow latency. Analysis occurs after Google has delivered the message and is asynchronous to mail flow, so no analysis latency of any stage is introduced into message delivery and Section 4.1.1.1.4 continues to hold. This is a structural advantage of the post-delivery architecture: a gateway holding mail in flight for detonation must choose between delaying delivery and releasing the message unanalyzed, and the proposed solution faces neither trade-off.

When an attachment is determined malicious — by intelligence, signature, structure, or detonation — removal executes across every affected mailbox within seconds of that determination, without administrator intervention, as required by Section 4.1.5.1. When Agency policy elects to remove a malicious attachment rather than the message, the offending part is stripped, an explained copy is inserted through the Gmail API, and the original is preserved encrypted and restorable.

4.1.3.4 Account Takeover (ATO) Monitoring

Account takeover monitoring draws on the Google Workspace Admin SDK Reports API and on mail behavior, correlating identity events with subsequent mailbox activity so that a compromise is identified by behavior rather than by user report. Admin reports read is not mail access. The following are monitored, as required by Sections 4.1.3 and 4.1.3.4:

- **Mailbox rule modification.** Creation and modification of forwarding, deletion, auto-move, and auto-archive rules — including the concealed auto-forwarding rules characteristic of an established takeover. The current active state of mailbox rules is captured for every enrolled mailbox during the retrospective pass described in Section 4.1.2.3, establishing the rule baseline against which subsequent changes are evaluated. A mailbox compromised before the State engaged may carry a concealed

forwarding rule that has been exfiltrating mail for months, and that rule survives password reset, session revocation, and message removal. Rules created by a compromised account are surfaced and removable under Section 4.1.5.1, and a standing report of all mailboxes with external forwarding is available to the State.

- **Anomalous login behavior.** New device, new geography, and unrecognized client evaluated against a learned per-user baseline rather than a fixed threshold.
- **Impossible travel.** Successive authentications from locations irreconcilable with elapsed time.
- **OAuth token misuse.** Third-party application grants to unrecognized or newly observed client identifiers, and scope escalation on existing grants.
- **MFA bypass attempts.** Enrolment of new authentication factors and changes to existing factors on an account, evaluated against that account's established authentication pattern.
- **Unauthorized API-scope grants.** Domain-level and user-level scope grants inconsistent with established administrative practice.

Identity signals are correlated with follow-on mail behavior, so an anomalous authentication followed by lateral phishing, rule creation, or unusual outbound volume escalates as a suspected takeover rather than being reported as two unrelated events. Findings reach the State's SOC through the SIEM and SOAR integration described in Section 3.11 with the evidence that produced them.

Protection Before Client Rendering and Downstream Ingestion

Zero-click and pre-render protection. Detection begins on a change notification published by Google at the moment of delivery, rather than on a polling interval, so analysis starts within seconds of a message landing in the mailbox and runs concurrently with — not after — the mail client's own synchronization. Where Agency policy elects quarantine, a malicious message is removed from the mailbox server-side, at the cloud layer, so a message determined malicious is withdrawn before a client that has not yet synchronized ever retrieves it, and is withdrawn from clients that have. Because remediation acts on the mailbox rather than on the endpoint, it reaches every device attached to that mailbox at once, with no endpoint agent and no device-side action. This is the mechanism by which content that would execute on render, preview, or automated processing is neutralized at the cloud inbox layer without any change to MX records or mail routing.

Collaborative and automated destinations. Because analysis and remediation operate at the mailbox through the Google Workspace API, protection extends to mail that is consumed by systems rather than read by a person. Group addresses, shared and delegated mailboxes, resource mailboxes, and service-account mailboxes are enrolled and analyzed on the same terms as user mailboxes, and a malicious message is removed from each mailbox that received it. Removal at the mailbox precedes onward handling: a message withdrawn from the mailbox is no longer available to be auto-forwarded to an external recipient or ingested by a downstream CRM, ticketing, or workflow integration reading that mailbox.

Automated destinations are therefore protected natively from within the cloud environment, with no rerouting of mail and no change to the State's integrations.

Adversarial content targeting AI assistants. Messages are inspected for content constructed to manipulate an AI assistant rather than to deceive a human reader: instruction-like text directed at an assistant, and text hidden from the reader but present to a parser — white-on-white and transparent text, zero-width and bidirectional control characters, off-screen and zero-size elements, and metadata-borne instruction payloads. Hidden content is extracted and evaluated as part of the message rather than discarded, both because it is the carrier for prompt-injection payloads and because the same concealment is used to hide the language signals that business-email-compromise detection depends on. Content of this kind is scored and dispositioned under Agency policy at the mailbox layer, before an assistant or copilot operating on the mailbox reaches it.

Outbound Detection and User-Reported Phishing

Outbound analysis. Mail sent from State mailboxes is analyzed on the same terms as inbound mail, satisfying the requirement that the platform detect malicious mail in either direction. Outbound analysis serves a different purpose from inbound: it is the primary evidence that a State account has been compromised. A mailbox that begins sending financial-request language to colleagues, sending to recipients outside its established correspondents, sending at volumes or hours inconsistent with its own history, or sending mail bearing the markers of a phishing campaign is surfaced as a suspected takeover, and correlated with the identity signals described in Section 4.1.3.4. Data-loss prevention is applied to outbound content, identifying sensitive information leaving the State by mail.

Because the platform is not in the delivery path, outbound coverage is detection and post-send response rather than pre-send blocking: a message identified as malicious or as carrying sensitive content is surfaced, and where Agency policy directs, remediated from the sender's mailbox and from the mailboxes of any internal recipients. Outbound remediation reaches State-managed mailboxes; it cannot recall a message already accepted by an external recipient's mail system, and the Vendor does not represent otherwise.

User-reported phishing. Reports submitted by State users enter the same analysis pipeline as automatically detected mail rather than a separate manual queue. A reported message is re-analyzed against current threat intelligence, and the platform identifies every other mailbox that received the same or a substantially similar message so the response covers the whole campaign rather than the single report. Where analysis confirms the report and Agency policy directs, removal executes across all affected mailboxes automatically. The reporting user receives an outcome notification, and the disposition — confirmed, benign, or inconclusive — is recorded for the reviewing analyst. Reports are treated as a detection signal in their own right: a cluster of reports naming related messages escalates ahead of any individual report.

4.1.5 Administrative and Incident Response Automation

4.1.5.1 Automated Remediation

When a message is determined malicious in any mailbox, the platform identifies every other mailbox that received it and removes it from all of them — across all domains, agencies, and organizational units, and including shared, delegated, resource, and service-account mailboxes — within seconds of that determination. Removal executes automatically under Agency policy, without manual administrator approval and without a per-message human decision, as required by Section 4.1.5.1.

The original message is retained encrypted with an integrity hash rather than destroyed. An administrator may review a quarantined message and release it back to the mailbox through an audited workflow, and the retained original is hash-verified before it is restored. Partial failure across a multi-mailbox action is compensated, so no mailbox is left in a half-remediated state.

Agency policy is configurable per detection family and per organizational unit, and governs which findings trigger automatic removal. The platform's default posture on installation is detect-only so that the State can observe detection quality before enabling automatic action; autonomous removal is enabled by the State per policy and, once enabled, requires no administrator involvement at the time of an incident.

SOAR-Triggerable Response Actions (Section 3.11)

The State's SOAR can invoke each of the remediation actions Section 3.11 requires, through an authenticated API:

- **Global message clawback.** Identify and remove a message from every mailbox that received it, across all domains, agencies, and organizational units, as described in Section 4.1.5.1.
- **Mailbox rule removal.** Disable or delete a malicious mailbox rule — a concealed external forwarding rule, or a rule auto-deleting security or finance correspondence. Prior rule state is captured before removal so the change is auditable and reversible.
- **Account-level threat response.** Terminate all active sessions for a compromised account, invalidating issued credentials so the attacker's existing access ends rather than persisting until token expiry, and notify the affected user out-of-band. Session termination, mailbox-rule removal, and message clawback are the three actions that together end an active takeover; performing any one alone leaves the compromise partially intact.

Each action is available to the State's SOAR and is equally available through the platform's native workflows, so the State is not required to operate a SOAR to obtain automated response. Every action is written to the hash-chained audit log with the acting principal, whether it was invoked by a human administrator, by a native workflow, or by the State's SOAR.

4.1.5.2 Role-Based Access Control

Administrative access is role-based and scoped, satisfying Section 4.1.5.2.

Agency separation. Each State agency is provisioned as a distinct organization within the platform. Every mailbox, detection, alert, quarantined message, and audit record carries the organization it belongs to, and that scope is enforced on every read. An administrator's visibility is bounded by the organizations they are assigned to, so department IT staff see logs, detections, and remediation history for their own agency and no other. Separation is enforced in the data-access layer rather than by interface filtering, so it holds across the web interface, the API, and exported records alike.

Roles. The platform distinguishes a State-wide administrator role, an agency-scoped administrator role, and standard user access. Agency-scoped administrators hold full administrative capability within their own agency — reviewing detections, releasing or purging quarantined mail, managing policy, and reading audit history — without visibility into any other agency. State-wide administration is reserved to explicitly designated personnel.

Delegated administration. Agency-scoped administration is delegated by the State: a State-wide administrator assigns administrators to an agency, and that assignment is itself an audited action. Delegation does not require vendor involvement, and the State may grant, change, or revoke it at any time.

Audit logging of administrative access. Every administrative action — policy change, quarantine release, message restore, role assignment, and access to mailbox-derived content — is written to a hash-chained audit log recording the acting principal, the affected object, the organization scope, and the time. Administrative reads are logged alongside administrative writes, so review of who accessed what is available to the State and not only to the vendor. Those records are exported to the State's SIEM on the same path as detection events, and are subject to the same tamper-evident retention.

4.1.6 Required Statement-of-Work Items

Section 4.1.6 of the solicitation requires vendors to provide answers to four items in their statement of work. The responses below address Section 4.1.6.1 (Architecture Verification Statement), Section 4.1.6.2 (Compliance Verification), and Section 4.1.6.3 (Deployment Timeline). Pricing under Section 4.1.6.4 is provided on the solicitation Pricing Pages and in the Cost Proposal section of this response.

4.1.6.1 Architecture Verification Statement

The proposed solution is an API-native advanced email-security layer for the State’s existing Google Workspace environment. It satisfies Section **4.1.1.1.1 No Gateway Routing**. No change to MX records, Workspace routing, or the Agency’s mail flow. The proposed solution is not in the delivery path. Email traffic flows natively into Google Workspace exactly as it does today.

Google remains the mail system of record, and delivery to the recipient mailbox is complete before any vendor processing begins.

Connection to the State's environment is established exclusively through Google Workspace APIs using domain-wide delegation, authorized by the State's Google Workspace super administrator and revocable by the State at any time. No software, agent, appliance, virtual machine, or credential is installed on any State-managed system.

The OAuth scopes requested are limited to those required to deliver the services described in this response, and are granted, scoped, and revoked by the State:

- **Mail read.** Retrieval of delivered messages and attachments for security analysis.
- **Mail modify.** Tag, trash, quarantine, surgical remediation, and restore of messages, as required for the automated remediation described in Section 4.1.5.1.
- **Directory read.** Enumeration of mailboxes across all domains and organizational units, for the automatic enrollment described in Section 4.1.6.3 and for impersonation and look-alike analysis. Directory read is not mail access.
- **Admin reports read.** Login and administrative event telemetry, for the account-takeover monitoring required by Section 4.1.3.4. Admin reports read is not mail access.
- **Mail settings read.** Enumeration of the current active state of mailbox rules, forwarding addresses, delegation, and auto-reply settings, for the mailbox-rule monitoring required by Section 4.1.3.4. Mail settings read does not grant access to message content.

No scope beyond those listed is requested. The scopes requested are:

- **Mail read —** https://www.googleapis.com/auth/gmail.readonly
- **Mail modify —** https://www.googleapis.com/auth/gmail.modify
- **Directory read —** https://www.googleapis.com/auth/admin.directory.user.readonly
- **Admin reports read —** https://www.googleapis.com/auth/admin.reports.audit.readonly
- **Mail settings read —** https://www.googleapis.com/auth/gmail.settings.basic

Mail flow

The following describes the complete path of a message under the architecture scored for Section 4.1.6.1. The State's MX records are unchanged at every step. Pub/Sub, directory, and reports are not part of this figure; they are shown separately under *Ingestion and directory* (Diagram B) so that reviewers do not read the proposed solution as sitting in mail flow.

Step	What happens	Vendor involvement
1	An external sender transmits to the State's published MX hosts, which resolve to Google.	None. MX records are unchanged and continue to point to Google.
2	Google Workspace applies its native security controls and delivers the message to the recipient mailbox.	None.
3	The message is in the mailbox and available to the user. Delivery is complete. It is not gated on any vendor system.	None.
4	After delivery, the message is read through the Gmail API.	API read, outside the delivery path.
5	Analysis is performed in Nelnet's FedRAMP-authorized environment: behavioral business-email-compromise and account-takeover detection, phishing and URL reputation, staged attachment analysis (hash intelligence, signature scan, static structural inspection, and dynamic detonation of what those stages do not resolve), and multi-engine data-loss prevention. No mail-flow latency is introduced.	Analysis only. The proposed solution is not in the delivery path.
6	Disposition is applied according to Agency policy — detect-only (the default; mailbox untouched), tag, trash, quarantine with the original retained encrypted in MnemoShare, or surgical remediation (insert a modified copy, verify it landed, then dispose the original). Restore is a first-class, auditable action. Findings are exported to the State's SIEM/SOAR and may drive native workflows.	Disposition and export, after delivery, through the Gmail API. Never SMTP re-injection.

Mail flow does not change; MnemoShare analyzes and remediates after Google has already delivered the message, typically quarantining a threat within seconds of delivery.

Detection Pipeline

Every message delivered to a monitored mailbox traverses the same pipeline, in the same order, whether it originated outside the State or between State mailboxes.

- **Notification.** Google publishes a change notification when a message is delivered. The platform pulls it from inside the hosting environment; nothing inbound is opened to receive it.
- **Retrieval and classification.** The message is retrieved through the Gmail API and classified as internal-to-internal or externally originated. That classification is derived from verified authentication results, not from the asserted sender domain, so an external message spoofing a State domain is not treated as internal.
- **Authentication evaluation.** Google's SPF, DKIM, and DMARC results are read as recorded at delivery and evaluated against the identity the message asserts.
- **Behavioral and language analysis.** The message is scored against the learned baseline for its sender-recipient relationship — cadence, established correspondents, typical send hours, dormancy — while a machine-learning classifier evaluates message intent, in particular financial-request intent. Identity analysis runs in parallel: display-name impersonation against the directory, look-alike and homoglyph domain comparison against known-good correspondents, reply-to divergence, and payment-detail change requests.
- **URL analysis.** URLs are extracted from body, headers, and attachments, including URLs encoded in QR images. Redirect chains are followed to the true final destination and scored against multiple reputation sources.
- **Attachment analysis.** Attachments enter the staged funnel described in Section 4.1.3.3 — hash intelligence, signature scanning, static structural inspection, and dynamic detonation for what remains unresolved.
- **Data-loss prevention.** Content is evaluated for sensitive information, using pattern matching and named-entity recognition as a second stage.
- **Scoring.** The signals above are combined into a single weighted judgment rather than applied as independent pass/fail rules. Weighting is configurable per detection family, and a signal that would be unremarkable alone escalates in combination — an unfamiliar sender is not itself notable, an unfamiliar sender making a payment-detail change from a look-alike domain is.
- **Correlation.** Mail findings are correlated with the identity telemetry described in Section 4.1.3.4, so an anomalous authentication followed by internal phishing escalates as a suspected takeover rather than as two unrelated events.
- **Disposition.** Agency policy, configurable per detection family and per organizational unit, determines the action, as described under Disposition, restore, and export.
- **Audit and export.** The finding, its supporting evidence, and any mailbox mutation are written to a hash-chained audit log and exported to the State's SIEM.

Behavioral, URL, attachment, and data-loss-prevention analysis execute concurrently rather than in series; a slow attachment verdict does not hold up a behavioral finding. Where a stage produces a verdict after disposition has already occurred — a detonation verdict, or a URL that turns malicious later — the finding re-enters at scoring and the resulting action executes on the same automated path, without administrator intervention.

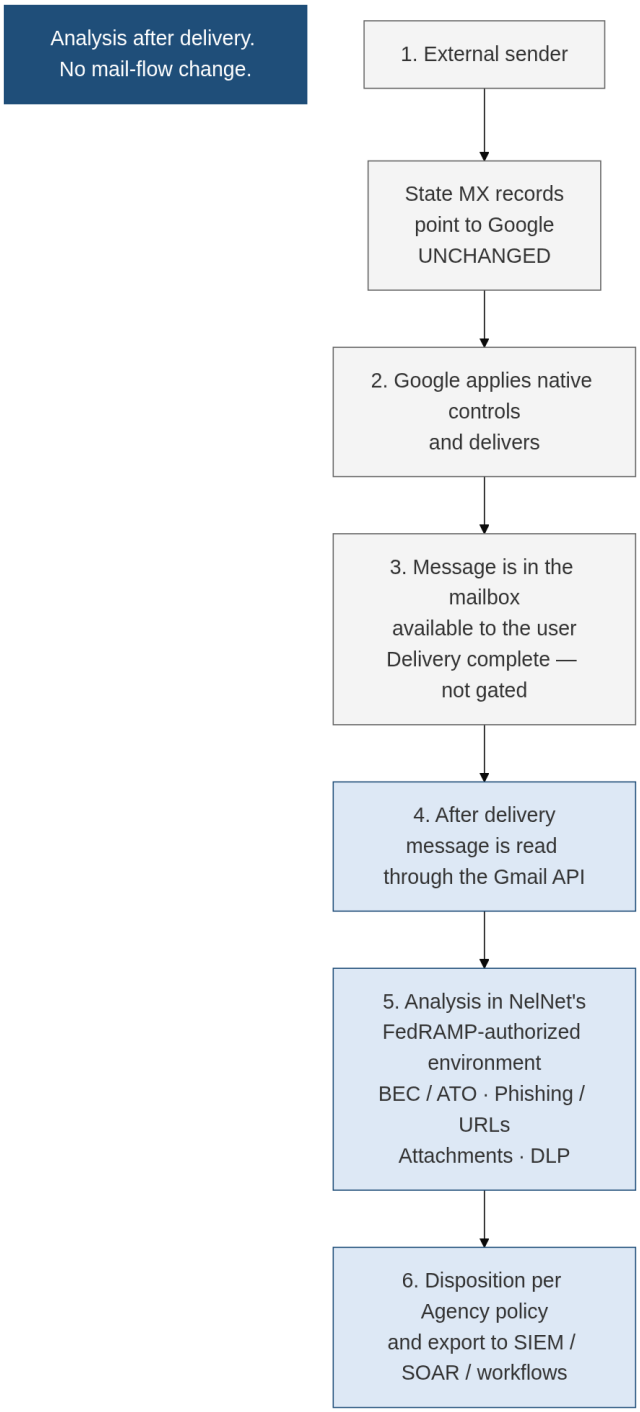

Figure A — Mail flow (scored for Section 4.1.6.1). Steps 1–3 are Google delivery. The proposed solution appears only after delivery.

Analysis performed after delivery

Hosted within Nelnet's FedRAMP-authorized environment. After the Gmail API read in step 4, the same detection stack runs on every retrieved message. Analysis does not add latency to mail delivery.

- **Business-email compromise and account takeover.** Comparable behavioral BEC and account-takeover detection, plus content remediation and secure transfer. Directory-backed identity modeling supports display-name impersonation, user-address and domain look-alikes, reply-to mismatches, and relationship signals. Coverage includes internal-to-internal mail that Google delivers inside the tenant and that never crosses the public internet. That visibility is native to the API design; no MX or routing change is required to see it.
- **Phishing and URLs.** URLs and registrable domains are extracted from the message and evaluated for reputation, and those verdicts are correlated with impersonation and authentication-results signals. Google's Authentication-Results (SPF / DKIM / DMARC) are consumed as inputs, not replaced.
- **Attachments.** Staged attachment analysis as described in Section 4.1.3.3: file-hash intelligence, signature scanning and static structural inspection performed in-boundary, with dynamic detonation applied to what those stages cannot resolve. Message content, message bodies, and data-loss-prevention analysis remain within the hosting environment.
- **Data-loss prevention and surgical content remediation.** Multi-engine DLP identifies sensitive content. When Agency policy elects remediation, the offending span or attachment is removed into authenticated MnemoShare storage, a modified copy is inserted through the Gmail API, the insert is verified, and the original is then disposed per policy. The original is preserved under audit so a false positive can be restored.

Disposition, restore, and export

Agency policy, configurable per detection family and per organizational unit, chooses what happens after a finding. The default is detect-only: alert and audit, with the mailbox left untouched. Stronger dispositions are Agency elections, not the factory default.

Agency-facing action	What happens	Can the user still open the message?
Detect-only (default)	Alert and audit. Mailbox untouched.	Yes. This is monitor mode.
Tag	A Gmail label is applied. The message stays in the mailbox.	Yes.
Remediate-only	A modified copy is inserted; the original is left in place.	Yes.

Trash	The message is moved to Gmail Trash. It remains user-recoverable in Gmail.	Yes, via Gmail Trash.
Quarantine	The original is removed from the mailbox and retained encrypted in MnemoShare, with an integrity hash. Administrators review and release.	No. This is the user-unreachable option. Quarantined within seconds of delivery.
Restore	The retained original is hash-verified and inserted back into the mailbox through an auditable workflow.	Restored.

When Agency policy elects quarantine or surgical remediation, the write occurs **after delivery** through the Gmail API. Surgical remediation always uses this order: insert a copy, verify the copy landed, then dispose the original. The reverse order is not used. Mail is never re-injected by SMTP.

Restore is first-class. Because the original is retained with an integrity hash, an administrator can return it to the mailbox through an auditable workflow. Partial failure is compensated so a mailbox is not left in a half-restored state.

Every detection and every mailbox mutation is written to a hash-chained audit log. Those records are exported to the State's SIEM or SOAR, and the proposed solution's native workflows may run restore, notify, and multi-step response. The State may use its SOAR, the proposed solution's workflows, or both.

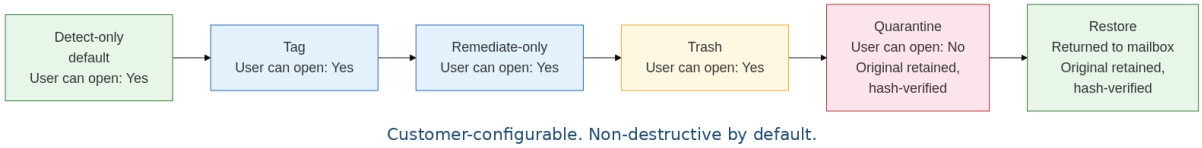

Figure C — Disposition options. Least invasive at left; restore returns a quarantined original to the mailbox.

Ingestion and directory

The following describes how the proposed solution learns that a message has been delivered and how it enrolls mailboxes. This is not mail delivery. Reviewers should not read this figure as a hop in the mail path.

- Each enrolled mailbox has a Gmail watch. Google publishes a change notification to Google Cloud Pub/Sub.

- MnemoShare **pulls** that subscription from inside Nelnet's AWS environment. Nothing inbound is opened into the enclave to receive mail or a webhook.
- The Admin SDK Directory API supplies enrollment and impersonation context. It is not mail access.
- The Admin SDK Reports API provides account takeover login and administrative events. It is not mail access.
- A scheduled reconcile covers any mailbox whose watch or history cursor has gone stale, so silent coverage loss does not persist.

Scale and Coverage Assurance

The deployment is sized for the full mailbox population identified in Section 2, including the shared, delegated, resource, and service-account mailboxes the State has confirmed are in scope. Analysis is distributed across horizontally scaled workers, so throughput is a function of allocated capacity rather than a fixed ceiling, and the population can grow without redesign.

Coverage is verified rather than assumed. Each enrolled mailbox holds an active change subscription, and a scheduled reconciliation identifies any mailbox whose subscription or processing cursor has gone stale and re-establishes it, so a silent loss of coverage on an individual mailbox is detected and corrected rather than persisting unnoticed. Mailboxes created after onboarding are enrolled automatically on the next directory synchronization. Coverage state is reportable to the State at any time, not only at onboarding.

Non-human mailboxes are enrolled and analyzed, and are modeled distinctly: a service account that sends automated notification mail at a constant rate has a legitimate behavioral profile that would be anomalous for a person, and is baselined on its own terms rather than against human sending patterns.

SIEM and SOAR Integration

Detection events, disposition actions, and administrative actions are exported to the State's SIEM as structured JSON records, on a delivery cycle as frequent as one minute, satisfying the near-real-time export required by Section 3.11. Records carry the finding, the evidence that produced it, the affected mailboxes, the action taken, and the acting principal, so an analyst can reconstruct a detection without leaving the SIEM.

Remediation does not depend on the State's SOAR. The platform executes automated response natively — quarantine, surgical remediation, restore, mailbox-rule removal, and multi-step response — through its own workflow engine, with no external playbook required and no per-message human approval step, as required by Section 4.1.5.1. The State's SOAR is an integration option, not a dependency.

Where the State elects to drive response from its own SOAR, integration is API-based in both directions: the platform exposes an authenticated API the State's SOAR can invoke to trigger remediation, and accepts webhook triggers issued by the State's SOAR, with outbound webhooks and lifecycle events available for the State's platform to subscribe to. The State may use its SOAR, the platform's native workflows, or both. The final integration pattern is determined collaboratively during implementation.

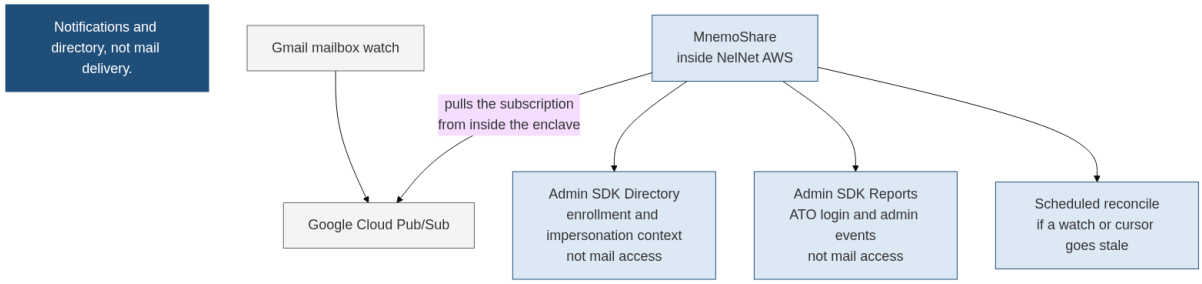

Figure B — Ingestion and directory. The proposed solution pulls notifications from inside Nelnet AWS. Directory and reports are not mail access. Do not place Pub/Sub on Diagram A.

Conformance with Section 4.1.1 (Core Architectural Requirements)

- **4.1.1.1.1 – No Gateway Routing.** No change to MX records, Workspace routing, or the Agency's mail flow. No inbound, outbound, or internal email is routed through a vendor appliance, mail relay, or cloud proxy. No journaling or SMTP relay configuration is used. The proposed solution is not in the delivery path.
- **4.1.1.1.2 – API-Based Deployment.** The platform connects directly to the State's Google Workspace environment using Google Workspace APIs, operating exclusively inside the perimeter as defined in Section 3.19.
- **4.1.1.1.3 – Zero Infrastructure Footprint.** The solution is delivered as cloud-native software as a service. It requires no on-premises hardware, no virtual appliances, no local agents, no customer-managed virtual machines or cloud compute resources, and no local software installation.
- **4.1.1.1.4 – Latency-Free Evaluation.** All security analysis occurs asynchronously and post-delivery. Because no vendor system is positioned in the delivery path, zero artificial latency is introduced into standard communication flows. This property is architectural rather than a performance target.

Failure Isolation and Reversibility

Because the solution carries no mail-routing dependency, it cannot become a single point of failure for State email. If the service were degraded or unavailable, mail delivery would continue normally and without interruption. Disconnection is equally non-disruptive: revoking the domain-wide delegation ends the integration immediately, with no cutover, no MX change, no maintenance window, and no effect on mail flow.

Behavior during a provider API interruption. Section 4.1.1.1.1 guarantees that mail delivery is unaffected by the platform's availability. The converse also holds: if the Google Workspace API is itself unavailable, mail continues to be delivered by Google and the platform resumes analysis when the API returns. Processing position is checkpointed per mailbox, so messages delivered during an interruption are analyzed on recovery rather than skipped — an outage delays analysis, it does not create an unexamined gap in coverage. The State is alerted when ingestion for a mailbox or for the tenant has stalled beyond a configured threshold, so a prolonged interruption is visible rather than silent.

Coexistence with Google Workspace and Gemini

In detect-only mode, mail is left untouched and message bodies remain as Google delivered them. When time-of-click URL protection is enabled, only link targets are re-pointed; message content, structure, and rendering are otherwise unchanged. In either mode, Gemini smart replies, AI-generated summaries, workspace context lookups, and native message rendering continue to operate on messages that remain in the mailbox, as required by Sections 4.1.2.2 and 3.17.

When Agency policy elects quarantine or surgical remediation, those actions occur after delivery via the Gmail API: the proposed solution inserts a copy, verifies that the copy landed, and then disposes of the original. Mail is never re-injected by SMTP. Native Workspace features, including Gemini, continue to operate on messages that remain in the mailbox. A message that policy has quarantined is no longer in the mailbox, by design; restore returns it through the same API path.

Section 4.1.5.1 requires removal or release of a message identified as malicious. The proposed solution performs that removal or release after delivery, under Agency policy, and not by sitting in the delivery path.

US Data Residency

All customer data — including email metadata, threat logs, and analysis artifacts — is stored and processed exclusively within data centers physically located in the United States, as required by Sections 4.1.4.1 and 3.8. The proposed solution is hosted within Nelnet’s environment at the FedRAMP Moderate baseline described under Section 4.1.6.2. No customer data is stored or processed outside the United States.

What is retained. Analysis reads mailbox content; it does not accumulate a copy of the State's mail. The platform retains the detection record and the evidence supporting it, derived relationship statistics and behavioral baselines, and the audit record of every action taken. Message bodies and attachment content are not retained beyond the analysis that produced a finding, with one exception: where Agency policy elects quarantine or surgical remediation, the original message is retained encrypted with an integrity hash precisely so it can be restored, and is disposed under the retention policy the State configures. Retention periods for detection records and audit history are configurable by the State and

default to the period required by the State's audit obligations. All retained data is stored within the United States as described above.

A formal architecture diagram illustrating the mail flow (Diagram A), ingestion and directory (Diagram B), and disposition options (Diagram C) is provided as Figures A, B, and C within this section.

4.1.6.2 Compliance Verification

The proposed solution satisfies Pathway A — Federal Cloud Alignment (FedRAMP) under Section 4.1.4.2.

The proposed solution is deployed and operated within Nelnet's FedRAMP-authorized AWS environment, at the FedRAMP Moderate baseline, under the Marketplace listing and Package ID identified below. The State is therefore served from an environment assessed against the FedRAMP Moderate baseline.

Supporting details required by Section 4.1.4.2.1 are provided below. The designation is current and listed on the FedRAMP Marketplace; it is not a FedRAMP Legacy status, which Section 4.1.4.2 excludes. Section 4.1.4.2 accepts Authorized, In-Process, and Ready designations.

Item	Detail
Marketplace listing name	AWS US East/West
Authorized entity	Amazon
Cloud Service Offering	AWS US East/West
Baseline	Class C (Moderate)
Designation	FedRAMP Certified
FedRAMP Package ID (PID)	AGENCYAMAZONEW
Relationship of proposed solution	Deployed within the authorization boundary of the Cloud Service Offering identified above.

All customer data is stored and processed exclusively within the United States, as required by Section 4.1.4.1 and as described in the US Data Residency statement above.

Data Isolation and Security Controls

The following controls are verified independently of the compliance pathway selected, and correspond to the Strict Data-Isolation Protocols enumerated in Sections 4.1.4.4.1 through 4.1.4.4.4:

- **Tenant isolation (4.1.4.4.1).** The State is served by a dedicated deployment provisioned for the State of West Virginia. State data is logically isolated from other tenants and is not co-mingled in shared application or data stores.
- **Encryption (4.1.4.4.2).** Data is encrypted at rest using AES-256 and in transit using TLS 1.3, which is enforced. Cryptographic operations use a FIPS 140-3 validated cryptographic module (Go official module, NIST CMVP).
- **Vendor access (4.1.4.4.3).** Vendor employees have no un-audited access to mailbox contents. Administrative access is role-based, individually attributed, logged, and auditable, and is granted only as required to deliver contracted support.
- **Framework alignment (4.1.4.4.4).** Controls are aligned to NIST SP 800-53. Under the FedRAMP pathway identified above, that alignment is evidenced by third-party assessment against the NIST SP 800-53 baseline rather than by vendor self-attestation.

4.1.6.3 Deployment Timeline

Onboarding and full-mailbox API verification are completed within 24 hours, without any change to mail routing and without disruption to active State communications.

The 24-hour period begins when the State's Google Workspace super administrator authorizes domain-wide delegation. That authorization is the single action that starts the deployment, and it is scheduled with the State in advance at a date and time of the State's choosing.

Phase	Elapsed	Activity	State effort
1	0:00 – 0:30	The State's Google Workspace super administrator authorizes domain-wide delegation for the agreed scope set. This is the only State action required to establish the connection.	One administrator, approximately 30 minutes.
2	0:30 – 4:00	All mailboxes are enrolled automatically by synchronizing the Google Workspace directory — across all domains, agencies, and organizational units, and including shared, delegated, and resource mailboxes. Mailbox watches are registered during enrollment as part of API verification. No mailbox list is manually maintained or supplied by the State.	None.
3	4:00 – 16:00	Per-mailbox API reachability is verified and recorded across the full enrolled population,	None.

		establishing provable coverage rather than assumed coverage.	
4	by 24:00	A coverage report is delivered to the State reconciling enrolled and verified mailboxes against the Workspace directory, and identifying any mailbox that could not be verified together with the reason.	Review and acceptance.

Non-Disruption

No phase of onboarding touches mail flow. There is no MX change, no cutover, no relay configuration, and no maintenance window. Mail delivery to all mailboxes continues without interruption throughout onboarding and at every point thereafter. Should the State elect to discontinue the service at any stage, revoking the domain-wide delegation ends the integration immediately and without effect on mail delivery.

Relationship between onboarding and historical scanning. The 24-hour commitment above covers onboarding and full-mailbox API verification, as Section 4.1.6.3 requires. Historical scanning is a distinct activity: it begins immediately upon activation, runs concurrently with and after enrollment, and completes on a schedule determined by the size of the agreed retrospective window and the tenant's available API quota. Retrospective processing of a ninety-day window across the full mailbox population is a substantially larger workload than onboarding, and the Vendor does not represent that it completes within the 24-hour onboarding window.

Two properties make that separation safe for the State. First, protection is not gated on it: live monitoring, detection, and remediation are fully operative on newly delivered mail from the moment a mailbox is verified, independent of how much history has been processed. Second, it consumes nothing the State must supply — no State effort, no infrastructure, and no mail-flow involvement at any point. A schedule for retrospective completion is agreed with the State at implementation and reported against on the coverage report described in Section 4.1.2.3.

Onboarding scales with directory size rather than with manual effort, so the 24-hour commitment holds at the approximately 22,000 mailboxes identified in Section 2 and on the Pricing Pages, and continues to hold as the State’s mailbox population changes. Mailboxes created after onboarding are enrolled automatically on the next directory synchronization.

Service Availability

The Vendor commits to 99.9% monthly availability of the detection and administrative service. Because analysis occurs after native delivery through the Google Workspace API and is not in the mail path, State mail delivery is unaffected by the availability of the detection service in all events: mail continues to be delivered by Google Workspace without interruption, and no message is delayed, queued, or lost as a result of a service interruption.

State Onboarding Responsibilities

The deployment requires no State infrastructure work, no network changes, no mail routing changes, and no software installations. The State inputs required during the scheduled 24-hour window are limited to the following, and are confirmed with the State before the window opens:

- **Administrator availability.** A Google Workspace super administrator available at the agreed start time to authorize domain-wide delegation.
- **Scope confirmation.** Confirmation of the domains and organizational units in scope for enrollment.
- **Technical point of contact.** A named State technical contact reachable during the window to resolve any mailbox that cannot be verified.
- **Coverage report acceptance.** Review and acceptance of the coverage report delivered at the close of the window.

Because the solution is not in the mail path, none of these items carries risk to State communications. If any is delayed, mail delivery is unaffected, and onboarding simply resumes when the State is ready.

Continuing State Responsibilities

Beyond the onboarding window, the State's continuing involvement is limited to the following. These describe the normal operation of the service and are not conditions on Vendor performance.

- **Google Workspace environment.** The State maintains its Google Workspace tenant and the domain-wide delegation authorization that permits API access. No additional administrative work is introduced.
- **Directory accuracy.** Mailbox enrollment follows the Google Workspace directory as the State maintains it today. No separate roster is kept, and no reconciliation activity is asked of the State between annual renewals.
- **Technical point of contact.** A named State contact remains reachable for coverage questions and configuration changes.
- **Disposition policy ownership.** The State designates the personnel authorised to approve quarantine, remediation, and release policy, and to approve changes to it.

4.1.6.4 Pricing

Pricing under Section 4.1.6.4 is provided on the solicitation Pricing Pages and in the Cost Proposal section of this response. Unit prices are not restated on this architecture page.

Platform Scope and Microsoft TOC

The solution proposed here is delivered natively against Google Workspace, which is the Agency's environment today. Microsoft 365 support is on the Vendor's product roadmap and is not offered as a delivered capability in this response. A Microsoft Graph adapter exists in the platform today; native Microsoft 365 functional parity — historical mailbox scanning, automated remediation, mailbox-rule monitoring, and account-takeover detection — is roadmap work and is described here as such rather than as delivered. No mail-routing change, journaling, or gateway-based routing would be introduced on either platform.

Section 5.0 — Palo Alto NGFW Compatibility

The proposed solution does not sit in the network data path and does not route, proxy, or inspect network traffic. All communication occurs between Google Workspace APIs and the vendor environment; no traffic traverses, terminates on, or depends upon the State's network perimeter. The solution therefore requires no change to the State's Palo Alto NGFW platform, introduces no architectural conflict with it, and does not interact with or rely upon the State's existing Palo Alto infrastructure. Consistent with the State's clarification, traditional NGFW compatibility testing is not applicable to an API-integrated solution of this architecture.

Tier-3 Support

The State is given direct access to the Vendor's advanced engineering support, not to a first-line helpdesk with escalation behind it. A named technical contact is assigned at implementation, and State staff reach engineering directly for issues arising in both implementation and ongoing operation — detection tuning and false-positive investigation, remediation and restore questions, SIEM and SOAR integration, API and scope questions, and analysis of a specific message or campaign. Support is included in the subscription with no per-incident or per-ticket charge, and no support tier is sold separately.

Exceptions and Clarifications

Clarification regarding Sections 4.1.1.1 and 4.1.3.2 (link handling). Real-time inspection at the moment of the click, as required by Section 4.1.3.2 and defined in Section 3.13, is achievable only where the activated link resolves through an inspection step at the time of the click. The proposed solution provides this server-side and post-delivery: the only change to a delivered message is the link target, which transparently preserves and displays the original destination. Mail is not routed through any gateway, relay, or proxy; MX records and mail flow are unchanged; and no software is installed on any State device. The original message is preserved intact and can be released to the mailbox unchanged at any time. This is the means by which the proposed solution delivers the click-moment inspection Section 4.1.3.2 requires while preserving the mail-flow architecture and device-footprint protections of Section 4.1.1.1.

Exception — Section 4.1.1.1.5 (Dual-Platform Flexibility). The solution proposed here is delivered natively against Google Workspace, which is the Agency's environment. Native Microsoft 365 functional parity is on the Vendor's product roadmap and is not offered as a delivered capability in this response. No mail-routing change, journaling, or gateway-based routing would be introduced on either platform.

4.1.6.4 Pricing Structure

Service	Cost
Line 1 Core Detection Core Detection per Mailbox estimated 22,000 Mailboxes	**$16.45 per mailbox / year** 22,000 x $16.45 = $361,900.00
Line 2 Core Detection - Year 2 (Optional Renewal) Core Detection per Mailbox estimated 22,000 Mailboxes	**$16.94 per mailbox / year** 22,000 x $16.94 = $372,680.00
Line 3 Core Detection - Year 3 (Optional Renewal) Core Detection per Mailbox estimated 22,000 Mailboxes	**$17.45 per mailbox / year** 22,000 x $17.45 = $383,900.00
Line 4 Core Detection - Year 4 (Optional Renewal) Core Detection per Mailbox estimated 22,000 Mailboxes	**$17.98 per mailbox / year** 22,000 x $17.98 = $395,560.00
Four-Year Total (Base Year + 3 Optional Annual Renewal Years)	**$1,514,040.00**

* The usage amount provided is an estimate only. Actual usage may be higher or lower, and the figure is not guaranteed.

Pricing Structure — Flat, No Metering

Pricing is a flat annual subscription at a single per-mailbox rate, with no additional or conditional charges of any kind. A firm per-mailbox rate is stated on the Pricing Pages for the base year and for each Optional Annual Renewal Year, and is carried to the Cost Proposal section of this response.

- **All updates and new releases.** Every software update and detection-model improvement is included at no additional charge. There is no separate maintenance or upgrade fee.
- **Tier-3 support.** Vendor tier-3 engineering support is included for the duration of the term.
- **Incident response automation.** The disposition and remediation actions described in Sections 4.1.3 and 4.1.5.1 — including quarantine, surgical remediation, restore, and mailbox-rule removal — are included, with no per-incident, per-action, or per-remediation charge.
- **The complete detection capability set.** Business email compromise and spoofing, advanced phishing and credential harvesting, malware and zero-day attachment inspection including dynamic detonation, and account-takeover monitoring are all included in the single rate. No capability is licensed as a separate module.
- **All analysis, including machine-learning classification.** Behavioral analysis and language-pattern classification are included with no per-message, per-scan, or per-analysis charge, and the State is not required to hold or fund any separate artificial-intelligence subscription.

- **SIEM and SOAR integration.** Threat-event export and SOAR-triggered response are included.
- **Onboarding and deployment.** The onboarding described in Section 4.1.6.3, including the coverage report, is included. There is no implementation, professional services, or activation fee.

Charges That Do Not Apply

The State is asked to note the following explicitly, because they are common in this product category and none of them applies here:

- No platform fee, subscription minimum, or minimum annual contract value.
- No per-module licensing. Capabilities are not sold separately or in tiers.
- No per-message, per-mailbox-scanned, per-analysis, or per-remediation metering, and therefore no overage charge under any level of mail volume.
- No charge tied to the volume of mail the State sends or receives. The subscription is priced based solely on mailbox count, so the State's cost does not vary with message traffic.

Term, Renewals, and Price Firmness

Pricing is provided for the base year and for Optional Annual Renewal Years 2, 3, and 4 as required by Section 6.0. A separate per-mailbox rate is quoted for each of the four years on the Pricing Pages, and each rate is firm for the full year to which it applies. No price adjustment provision of any kind applies: there is no escalator, no indexation to any published index, no surcharge, and no renewal uplift. Optional Annual Renewals are initiated by the State, agreed to by the Vendor, and executed via Change Order.

The quantity on the Pricing Pages is an estimate, and the subscription is calculated as the per-mailbox rate multiplied by the actual mailbox count. Because a portion of the cost of delivering the service — tier-3 support, continuous monitoring participation, and operation of the deployment itself — is incurred per deployment rather than per mailbox, the quoted rate is based on the approximately 22,000 mailboxes identified in Section 2. The following applies:

- **Growth.** Mailboxes added at any point during a term are billed at the per-mailbox rate in effect for that term for the remainder of that term. There is no mid-term uplift, no re-tiering, and no minimum increment.
- **Reduction.** The per-mailbox rate for a given contract year holds without change at any mailbox count of 17,600 or above, which is eighty percent of the estimated quantity. Should the population fall below that level, the parties review and agree the per-mailbox rate for the affected period and execute it via Change Order, using the same mechanism Section 6.0 provides for Optional Annual Renewals.
- **Reconciliation.** Mailbox count is reconciled against the Google Workspace directory annually at renewal rather than continuously, so the Agency’s invoice does not fluctuate through the year.

Notes

Customer Responsibility. The engagement requires no State infrastructure work, no network or mail-routing changes, and no software installation on any State device. State involvement is limited to the onboarding inputs described in Section 4.1.6.3 and the continuing items listed there. A single named State technical point of contact coordinates scheduling and any subsequent configuration changes. No fee of any kind is charged for scheduling, rescheduling, or configuration changes at any point during the term.

Third Party Review and Engagement. Any direct engagement with a Third-Party vendor for CampusGuard services will require a separate agreement with commercially acceptable terms and pricing established between the Third-Party vendor and CampusGuard.

Limitations of Liability. IN NO EVENT WILL CUSTOMER ON THE ONE HAND, OR CAMPUSGUARD AND/OR ITS LICENSORS AND CONTRACTORS, ON THE OTHER HAND, BE LIABLE TO THE OTHER, WHETHER IN CONTRACT, TORT OR ANY OTHER LEGAL THEORY (INCLUDING, WITHOUT LIMITATION, STRICT LIABILITY AND NEGLIGENCE), FOR LOST PROFITS OR REVENUES, LOSS OF USE OR LOSS OF DATA, OR FOR ANY INDIRECT, SPECIAL, EXEMPLARY, PUNITIVE, MULTIPLE, INCIDENTAL, CONSEQUENTIAL, OR SIMILAR DAMAGES ARISING OUT OF OR IN CONNECTION WITH THE SERVICES OR DELIVERABLES, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT WILL THE LIABILITY OF CAMPUSGUARD AND ALL OF ITS EMPLOYEES, AGENTS, LICENSORS AND CONTRACTORS EXCEED, IN THE AGGREGATE, AN AMOUNT EQUAL TO THE TOTAL AMOUNT OF FEES ACTUALLY PAID BY CUSTOMER TO CAMPUSGUARD UNDER THE APPLICABLE SERVICE ORDER UNTIL THE DATE CUSTOMER SUBMITS ITS INITIAL CLAIM FOR SUCH LIABILITY. NO ACTION REGARDING SERVICES OR DELIVERABLES MAY BE BROUGHT AGAINST CONSULTANT MORE THAN TWO (2) YEARS AFTER THE CAUSE OF ACTION ACCRUES.