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: 000000180526

Legal Name: SUN MANAGEMENT INC

Alias/DBA:

Total Bid: $1,181,400.00

Response Date: 08/17/2026

Response Time: 17:27

Responded By User ID: linusroman

First Name: Linus

Last Name: Roman

Email: linus@sunmanagement.net

Phone: 804-690-7399

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 ESR08172600000001070	1

VENDOR

000000180526
SUN MANAGEMENT INC

Solicitation Number: CRFQ 0231 OOT2700000005

Total Bid: 1181400 **Response Date:** 2026-08-17 **Response Time:** 17:27:34

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	12.400000	272800.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments: We are quoting the following Abnormal Security SKU's...
#IES "Harness advanced behavioral AI to block socially-engineered attacks and other malicious emails. (1 Year Subscription)
Abnormal AI Inc - IES
#CORE-ATO "Core Account Takeover Protection
Abnormal AI Inc - CORE-ATO
#AISM "AI Security Mailbox
Abnormal AI Inc - AISM
#PFF "Platform Fee
Abnormal AI Inc - PFF

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	13.050000	287100.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments: We are quoting the following Abnormal Security SKU's...
#IES "Harness advanced behavioral AI to block socially-engineered attacks and other malicious emails. (1 Year Subscription)
Abnormal AI Inc - IES
#CORE-ATO "Core Account Takeover Protection
Abnormal AI Inc - CORE-ATO
#AISM "AI Security Mailbox
Abnormal AI Inc - AISM
#PFF "Platform Fee
Abnormal AI Inc - PFF

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	13.750000	302500.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments: We are quoting the following Abnormal Security SKU's...
#IES "Harness advanced behavioral AI to block socially-engineered attacks and other malicious emails. (1 Year Subscription)
Abnormal AI Inc - IES
#CORE-ATO "Core Account Takeover Protection
Abnormal AI Inc - CORE-ATO
#AISM "AI Security Mailbox
Abnormal AI Inc - AISM
#PFF "Platform Fee
Abnormal AI Inc - PFF

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	14.500000	319000.00

Comm Code	Manufacturer	Specification	Model #
43230000			

Commodity Line Comments: We are quoting the following Abnormal Security SKU's...
#IES "Harness advanced behavioral AI to block socially-engineered attacks and other malicious emails. (1 Year Subscription)
Abnormal AI Inc - IES
#CORE-ATO "Core Account Takeover Protection
Abnormal AI Inc - CORE-ATO
#AISM "AI Security Mailbox
Abnormal AI Inc - AISM
#PFF "Platform Fee
Abnormal AI Inc - PFF

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.

Requirement (RFQ CRFQ OOT2700000005 — Section 4. General Requirements)
4. GENERAL REQUIREMENTS
4.1 Contract Items and Mandatory Requirements
Vendor shall provide Agency with the Contract Items listed below on an open-end and continuing basis. Contract Items must meet or exceed the mandatory requirements as shown below.
4.1.1 Core Architectural Requirements (Non-Gateway)
(Context) To preserve native email delivery speeds, prevent MX record modifications, and eliminate single-point-of-failure routing risks, the State requires a native, API-integrated deployment model.
4.1.1.1.1 No Gateway Routing: The solution must not require changing the State's MX (Mail Exchanger) records or routing inbound, outbound, or internal email traffic through an external secure email gateway (SEG) appliance, mail relay, or cloud proxy.
4.1.1.1.2 API-Based Deployment: The proposed platform must connect directly to the State's Google Workspace environment using Google Workspace APIs.
4.1.1.1.3 Zero Infrastructure Footprint: The solution must be 100% cloud-native (SaaS) and require no on-premises hardware, virtual appliances, must not require customer-managed virtual machines or cloud compute resources in any form, or local software installations.
4.1.1.1.4 Latency-Free Evaluation: Security analysis must occur asynchronously or inline via API hooks post-delivery, ensuring there is zero artificial latency introduced to standard communication flows.
4.1.1.1.5 Dual-Platform Flexibility: While optimized immediately for Google Workspace/Gemini, the solution must provide full, native API-level functional parity across Google Workspace and Microsoft 365. Microsoft 365 support must include historical mailbox scanning, automated remediation, mailbox rule monitoring, time-of-click URL inspection, account-takeover detection, and access to all required Microsoft Graph API scopes. The solution must not require MX changes, journaling, SMTP relays, or gateway-based routing for either platform.
4.1.2 Gemini & Workspace Functional Integration
(Context) The solution must operate exclusively through Google Workspace APIs ('inside the perimeter') without altering mail routing or MX records. Internal-to-internal email traffic includes all messages exchanged between State-managed mailboxes across all domains and OUs. The solution must not disable, degrade, or interfere with Gemini features including smart replies, AI-generated summaries, context lookups, or native message rendering

4.1.2.1 Internal Threat Detection (East-West Traffic): Because it sits inside the API tier, the solution must analyze internal-to-internal state emails (employee to employee) for lateral phishing, data exfiltration, or account compromise. Internal-to-internal email traffic includes all messages exchanged between State-managed mailboxes, regardless of domain, OU, or routing path.

4.1.2.2 Coexistence with Gemini: The tool must not interfere with Gemini's native capabilities, such as AI-driven draft summaries, smart replies, or workspace context lookups, or disable, degrade, or interfere with Gemini features including smart replies, AI-generated summaries, context lookups, and native message rendering.

4.1.2.3 Historical Scanning: The API integration must allow the tool to scan existing historical mailboxes retrospectively to establish a baseline threat profile immediately upon activation. Historical scanning must include all existing mailbox contents, including messages, attachments, and mailbox rules, for the full retention period available within Google Workspace. Scanning must begin immediately upon activation and complete without requiring mail routing changes.

4.1.3 Technical & Threat Protection Capabilities

(Context) The solution must utilize machine learning that includes behavioral baselining, anomaly detection, and language-pattern analysis beyond static rule-based filtering. BEC detection must include internal and external impersonation, vendor compromise, invoice fraud, and look-alike domain analysis. Time-of-click URL inspection must analyze redirect chains, embedded URLs, QR-encoded links, and dynamically generated phishing pages. Malware and zero-day protection must include deep file inspection and cloud-based sandboxing for executables, archives, scripts, Office documents, PDFs, and containerized payloads. ATO monitoring must include mailbox rule changes, anomalous login behavior, impossible travel detection, OAuth token misuse, MFA bypass attempts, and unauthorized API-scope grants.

4.1.3.1 Business Email Compromise (BEC) & Spoofing: Must analyze language sentiment, communication baselines, and look-alike domains to stop executive impersonation, vendor invoice fraud, and employee payroll diversion tactics.
4.1.3.2 Advanced Phishing & Credential Harvesting: Must perform real-time URL sandboxing and dynamic link inspection at the time-of-click to block morphing malicious links.

4.1.3.3 Malware & Zero-Day Attachments: Must include deep file inspection and sandboxing capabilities for modern vector payloads.

4.1.3.4 Account Takeover (ATO) Monitoring: Must monitor mailbox rule modifications (e.g., sneaky auto-forwarding rules created by attackers) and anomalous login behaviors across the state infrastructure.

4.1.4 Compliance, Certifications, & Data Security

4.1.4.1 To maximize vendor competition while ensuring rigorous data security for State infrastructure, the vendor solution must meet one of the following compliance pathways. All customer data must be stored and processed exclusively within the United States.
4.1.4.2 Pathway A — Federal Cloud Alignment (FedRAMP): The vendor must hold a current FedRAMP Authorized (Moderate or High baseline), FedRAMP In-Process, or FedRAMP Ready designation listed on the official FedRAMP Marketplace. FedRAMP Legacy status is not acceptable.
4.1.4.2.1 Vendor should provide the FedRAMP Package ID (PID) for verification.
4.1.4.3 Pathway B (State Cloud Alignment): Hold a current StateRAMP Category 2 or 3 Authorization, or maintain a reciprocal, active state-level cloud authorization (e.g., TX-RAMP).
4.1.4.4 Pathway C (Commercial Enterprise Security Framework): Hold a validated SOC 2 Type II certification plus a signed addendum outlining Strict Data-Isolation Protocols. This protocol must explicitly verify the four sub-items below.
4.1.4.4.1 State of West Virginia data is logically or physically isolated from other commercial tenants.
4.1.4.4.2 The solution uses end-to-end data encryption both at rest (AES-256) and in transit (TLS 1.3).
4.1.4.4.3 Vendor employees have no zero-trust or un-audited access to mailbox contents.
4.1.4.4.4 Alignment with NIST SP 800-53 or ISO/IEC 27001 cybersecurity frameworks.
4.1.5 Administrative & Incident Response Automation

4.1.5.1 Automated Remediation (Search & Destroy): When a threat is detected in one mailbox, the tool must have API permissions to automatically find and claw back (quarantine) identical malicious emails across all 22,000 state mailboxes globally within seconds of detection. This must include messages in all domains, agencies, OUs, shared mailboxes, delegated mailboxes, and any mailbox where the malicious message exists. Automated remediation must execute without manual administrator approval and must support SOAR-triggered workflows.

4.1.5.2 Role-Based Access Control (RBAC): Must support granular administrative controls allowing specific state department IT staff access to view logs only for their own respective agencies. RBAC must support multi-agency separation, delegated administration, and audit logging of all administrative access.

4.1.6 Vendor Submission Requirements

(Context) Vendors responding to this RFQ must explicitly provide answers to the following in their statement of work:

4.1.6.1 Architecture Verification Statement: Provide an architectural narrative or diagram proving that email traffic flows natively into Google Workspace without passing through an intermediary MX-record route or gateway proxy. US data residency: All customer data must be stored and processed within the United States.
4.1.6.2 Compliance Verification: Explicitly state which verification pathway from Section 4 the platform satisfies, providing supporting documentation (e.g., FedRAMP ID number, StateRAMP dashboard link, or current SOC 2 Type II report).

4.1.6.3 Deployment Timeline: Confirm that onboarding and full-mailbox API verification can be achieved in under 24 hours without disrupting active state communications.
4.1.6.4 Pricing Structure: Provide flat SaaS pricing based on an annual subscription model for approximately 22,000 mailboxes, inclusive of all updates, tier-3 support, and incident response automation.

Legend: Rows highlighted in light orange had no matching response in Abnormal's prior CRFQ OOT2600000032 submission (drafted response for this RFQ (OOT2700000005). All other Column B responses were reused/adapted from that prior submiss new requirement is broader, e.g. Microsoft 365 feature parity, compliance pathway, or pricing) before final submission.

Vendor Response
Abnormal integrates natively via API with both Microsoft 365 and Google Workspace tenants. Deployment requires no MX record changes, no mail rerouting, and no journaling.
Abnormal integrates natively via API with Google Workspace tenants. Deployment requires no MX record changes, no mail rerouting, and no journaling.
Abnormal integrates natively via API with Google Workspace tenants. Deployment requires no MX record changes, no mail rerouting, and no journaling.
Abnormal is cloud-native and delivered as a cloud-based service (SaaS) solution. Abnormal does not offer an on-premise solution and requires no customer-managed infrastructure or hardware. Abnormal also does not require any end-user-installed agents, plugins, or clients.
Abnormal's API-first architecture requires no rerouting and no gateway bottleneck. The platform analyzes messages after delivery and automatically remediates threats by quarantining malicious emails. There is no latency impact on legitimate email flow.
Abnormal integrates natively via API with both Microsoft 365 and Google Workspace tenants and supports functional parity for historical mailbox scanning, automated remediation, mailbox rule monitoring, and account takeover detection. URL rewriting currently is only available for Microsoft 365. Deployment requires no MX record changes, no mail rerouting, no SMTP relays, and no journaling.
Abnormal has visibility into internal email communication (east-west traffic) and prevents lateral phishing by analyzing, judging, and remediating malicious emails sent from internal accounts, regardless of domain. OU

Abnormal's unique behavioral AI engine analyzes thousands of unique human-related signals, learning what's normal for your organization. The solution applies this understanding to detect deviations in sign-in activity, mail filter rules, security events, SaaS activity, and internal emails to detect employee-compromised accounts. The solution will build a timeline of activity events and can automatically remediate compromised accounts. Additionally, unlike traditional SEGs, the solution has visibility on internal email communication (east-west traffic). This allows Abnormal to prevent lateral phishing attacks, data exfiltration, and account compromise regardless of domain, OU, or routing path by analyzing, judging, and remediating malicious emails from internal accounts.
Abnormal integrates with Google Workspace via API and does not sit in the delivery path where Gemini's native features operate. Abnormal's actions are limited to post-verdict remediation (e.g., moving a message judged malicious) via API call after Google's own processing completes, which does not modify, block, or degrade Gemini's AI-driven capabilities such as draft summaries, smart replies, or workspace context lookups for legitimate mail.
Abnormal analyzes 30 to 90 days of historic emails to build initial detection baselines with no lift from the customer. Abnormal’s API-based install is completed in minutes and requires no MX changes, no mail rerouting, and no journaling.
Abnormal's advanced AI/ML engine analyzes tens of thousands of signals, including content tone and urgency, interaction patterns, user titles, and authentication activity, to baseline known good behavior and detect advanced attacks like business email compromise, account takeover, and more. This approach enables proactive and adaptive threat detection that does not depend on predefined rules or

Abnormal uses AI to analyze human behavior and email content, analyzing thousands of unique human-related signals to learn what is normal for your organization. By leveraging the organization's past communication, it detects malicious actors and flags suspicious behavior like unexpected financial requests, even from seemingly legitimate senders, including VIP impersonation, invoice and payroll fraud, multi-staged credential phishing, and business process fraud. Abnormal detects vendor fraud using domain intelligence (domain age, lookalike analysis), behavioral modeling of vendor relationships, and topic classification models (invoice, billing, RFQ). Abnormal's Knowledge Hub also monitors historical vendor-to-employee interactions and aggregates risk levels across the Abnormal community to surface high-risk vendors and supply-chain threats.

Abnormal takes an AI-first approach to detecting malicious emails rather than relying on traditional sandboxing technologies. While conventional sandboxing typically depends on detonating files or comparing them against known threat intelligence, Abnormal evaluates emails using behavioral AI that understands normal communication patterns within your organization. Attachments, URLs, and other embedded content are assessed as part of a broader set of contextual and behavioral signals to determine whether an email poses a risk.

To identify threats, Abnormal examines attachments and links using static analysis, behavioral modeling, contextual intelligence, and content-based inspection within a secure, isolated environment. The platform analyzes common file types for indicators such as malware, malicious macros, executables, embedded URLs, QR codes, and other suspicious characteristics. URLs are also evaluated through reputation and content analysis, with optional URL rewriting available for additional time-of-click protection. This layered, AI-driven approach enables Abnormal to detect both known and previously unseen threats without depending on always-on dynamic sandbox detonation.

Abnormal takes an AI-first approach to detecting malicious emails rather than relying on traditional sandboxing technologies. While conventional sandboxing typically depends on detonating files or comparing them against known threat intelligence, Abnormal evaluates emails using behavioral AI that understands normal communication patterns within your organization. Attachments, URLs, and other embedded content are assessed as part of a broader set of contextual and behavioral signals to determine whether an email poses a risk. To identify threats, Abnormal examines attachments and links using static analysis, behavioral modeling, contextual intelligence, and content-based inspection within a secure, isolated environment. The platform analyzes common file types for indicators such as malware, malicious macros, executables, embedded URLs, QR codes, and other suspicious characteristics. URLs are also evaluated through reputation and content analysis, with optional URL rewriting available for additional protection. This layered, AI-driven approach enables Abnormal to detect both known and previously unseen threats without depending on always-on dynamic sandbox detonation.
Abnormal correlates authentication signals across email and SaaS apps (locations, devices, browsers, usage patterns) with external risk indicators from integrated applications (e.g., Google) to detect suspicious authentication attempts, including first-time VPN usage, unusual patterns of failed passwords, and hidden mail filter creation. Abnormal can immediately and automatically contain threats by signing the user out, disabling the account, and forcing a password reset.

United States Customer Data is stored only in the US. Abnormal uses subprocessors to host our data in cloud hosting environments. These subprocessors process (store) data on our behalf, but none of their personnel have access to the data. Select Abnormal personnel have access to Customer Data to perform service operations (e.g. troubleshooting / bug fix / support). While the data is hosted and processed in the United States, this personnel may have access to the data from their location, which may be outside of the United States. The locations where this personnel sit align with the country locations in the Affiliates section of our subprocessor list at www.abnormal.ai/subprocessors.
while Abnormal is adhering to Pathway C for this response, Abnormal has received FedRAMP Moderate Authorization with FedRAMP High in progress.
N/A
N/A
Abnormal is assessed by an independent auditor at minimum annually for SOC 2 Type 2 audit. Audit findings are assigned to an owner, tracked to remediation, and reported to management. See audit related documentation and reports at https://security.abnormal.ai.
Customer Data is logically segregated within our storage architecture using a unique tenant identifier.
Data in transit is encrypted using standard TLS 1.2 or higher with a 2048-bit RSA asymmetric key. HTTPS is required for all traffic. Data at rest is encrypted using 256-bit Advanced Encryption Standard (AES-256).
Abnormal protects customer data through a layered security program that includes AES-256 encryption at rest, TLS 1.2+ encryption in transit, role-based access controls, multi-factor authentication, and logical tenant isolation. Access to customer data is restricted by default and governed by a least-privilege model enforced through single sign-on. Access to malicious mailbox data is audited as well.
Customers can view our ISO27001:2022 certification in our Security Hub at https://security.abnormal.ai.

Beyond Abnormal's front-end, Inbound Email Security (IES) detection and auto-remediation controls at scale, Abnormal also addresses this automated remediation use case through two complementary, API-driven mechanisms:

Campaign Remediation, part of Abnormal's AI Security Mailbox, automatically clusters and removes every identical or near-identical copy of a confirmed malicious or spam message across the tenant's full mailbox population once a copy is confirmed — no per-message administrator sign-off required, provided the tenant is in Active Protection mode.

For SOC- and SOAR-driven response, Abnormal's Search & Respond exposes a REST API (POST /v1/search/remediate) that lets a SOAR platform search and bulk-remediate matching messages — up to 200,000 in a single action — across the entire mailbox population without requiring a human to log into the portal.

Additionally, every detection and remediation event, across email and other channels, surfaces in the unified Threat Log, which feeds SIEM/SOAR platforms. Abnormal maintains a native Cortex XSOAR content pack (currently v2.4.7) supporting threat/case retrieval, employee identity enrichment, and remediation commands, so XSOAR playbooks can query the Threat Log and trigger remediation automatically as part of an incident-response workflow.

Abnormal's platform supports Role-Based Access Control (RBAC) with default roles (Portal Global Admin, Portal Global Per-Product Admin, Portal Tenant Admin, Portal Tenant Per-Product Admin, Notification Recipient) plus customization options including granular read-only/read-write access, email content access levels, tenant-specific access, custom roles, and an audit log tracking user activities for compliance.

Please see our Network and Data Flow Diagram documentation at https://security.abnormal.ai. United States Customer Data is stored only in the US. Abnormal uses subprocessors to host our data in cloud hosting environments. These subprocessors process (store) data on our behalf, but none of their personnel have access to the data. Select Abnormal personnel have access to Customer Data to perform service operations (e.g. troubleshooting / bug fix / support). While the data is hosted and processed in the United States, this personnel may have access to the data from their location, which may be outside of the United States. The locations where this personnel sit align with the country locations in the Affiliates section of our subprocessor list at www.abnormal.ai/subprocessors.
Pathway C. Please see our Security Hub at https://security.abnormal.ai for our SOC 2 Type II report).

The average deployment timeline can be broken into several components, beginning with the initial setup, which takes less than a few minutes of customer effort. Full activation typically completes in ~5 days, depending on mail volume and organization size. The steps to deploy are:

- API Integration / Full Mailbox Verification (Takes about 5 minutes and does not disrupt active state communications): This is the complete Google Workspace solution integration with minimal configuration, including no MX record changes.
- Model Normalization (2 to 7 Days) Abnormal analyzes 30 to 90 days of historic emails to baseline behavior with no lift from the customer.
- Portal Access and RBAC Admins are granted portal access. Role-based access controls can then be configured for additional users.
- Organizational Rollout Planning Abnormal collaborates with the customer to align rollout timelines and provide relevant enablement and technical support materials.
- Configure Additional Modules (if applicable) Enable additional Abnormal add-on products and features (e.g., Account Takeover Protection, AI Security Mailbox) with permissions.
- Transition to Auto-Remediation In production mode, Abnormal begins fully active monitoring and automated remediation of malicious messages.

Given that the State of West Virginia - Office of Technology has already integrated with Abnormal in a Read-Only Capacity, these deployment items can be greatly accelerated since most of these steps have

Abnormal offers a subscription model with a starting price based on the number of protected mailboxes, with options for add-on features based on customer needs. Abnormal license entitlements match that of the Google Workspace licenses that Abnormal is protecting.

'Advanced Email Security, OT26077) and need a newly
sion and should be reviewed for exact fit (especially where the