Proofpoint AI Model Card

Threat Protection Group (TPG) Detection Systems

Version: 1.0 | Date: June 2026

---

1. Model Overview

Field	Detail
Product group	Threat Protection Group (TPG)
Capability areas	Email threat detection, content analysis, URL and web content analysis
Model owner	Proofpoint Threat Protection engineering leadership
Executive oversight	Chief AI & Data Officer; Assistant General Counsel

---

2. Intended Use

Primary purpose: Detection and prevention of cybersecurity threats — including malware, ransomware, phishing, business email compromise, and spam — across Proofpoint's Threat Protection product portfolio.

Target users: Enterprise and organizational customers using Proofpoint email and collaboration security products.

Out-of-scope uses: Proofpoint's detection systems are designed and validated exclusively for cybersecurity threat detection. Use outside this scope is not supported.

---

3. Model Capabilities

Proofpoint's Threat Protection detection systems analyze signals across email metadata, message content, and URLs and web content to identify and block cybersecurity threats. Detection systems operate across complementary capability areas that together provide multilayer protection.

---

4. Training Data

Proofpoint's detection models are trained on a combination of:

- Proofpoint threat intelligence data
- Labeled malicious message and URL samples
- URL and web content data
- Behavioral signals including sender and URL reputation

Customer data is used in model training only where necessary to provide and improve detection services, and only with appropriate controls and approvals in place.

---

5. Training Methodology

Not publicly disclosed in order to protect proprietary detection methods.

---

6. Evaluation

Models are evaluated prior to deployment using standard machine learning evaluation methodology, including precision-recall analysis against labeled datasets. Models are promoted to production only when they demonstrably outperform the incumbent model against defined quality thresholds.

---

7. Validation Methodology

All models undergo a shadow deployment validation process before production release:

- **Shadow deployment:** New models process live production traffic alongside the current production model without affecting outcomes, enabling direct performance comparison under real-world conditions.
- **Promotion criteria:** Models are promoted to production only after meeting defined performance thresholds evaluated against the incumbent model and known labeled datasets.
- **Rollback capability:** Model versioning is maintained across all detection systems, enabling rollback to previous versions where needed.
- **Change management:** New model versions are deployed side-by-side with existing models. Previous versions are retired only after the new version has demonstrated sustained performance.

---

8. Bias & Fairness

Part of Proofpoint's ISO 42001 certification.

---

9. Safety & Risk Evaluation

Part of Proofpoint's ISO 42001 certification.

---

10. Governance & Oversight

Governance Structure

Proofpoint maintains a defined AI governance structure with clear ownership of model risk decisions. Model risk authority is assigned to designated technical leads and engineering owners within each product division, with executive oversight from the Chief AI & Data Officer and Assistant General Counsel.

Regulatory Alignment

Proofpoint's AI governance program is informed by leading frameworks:

Framework	**Status**
ISO/IEC 42001:2023 — AI Management Systems	https://www.proofpoint.com/us/legal/trust/product-certifications

Framework	Status
NIST AI Risk Management Framework (AI RMF)	Informing governance program
SR 11-7 — Supervisory Guidance on Model Risk Management	Informing governance program
EU AI Act	Proofpoint's use is Low Risk

---

11. Ongoing Monitoring

Model performance is continuously monitored through:

- Volume tracking to ensure continued effectiveness at scale
- Continuous evaluation of model outputs against analyst-confirmed false positives and false negatives
- Retraining cycles incorporating production feedback to address evolving threat landscapes

Continuous Improvement Initiatives

Proofpoint is actively investing in the maturity of its AI governance practices, including:

- **ISO 42001 certification** — comprehensive AI management system implementation
- **Automated data quality assessment** — observability metrics and anomaly detection for training data
- **Automated pre-deployment validation** — pipeline planned for 2026
- **Data lineage and provenance** — automated tracking of data usage and provenance

---

12. Data Handling

Data Elements Processed

Proofpoint's detection systems process customer data as necessary to provide threat detection services. Data elements processed include:

- Email metadata (sender, recipient, subject line)
- Message content (indicators of threat, fraud, spam)
- URLs and associated web content

- Behavioral indicators (sender and URL reputation, delivery signals)
- Threat intelligence data

The specific data elements used vary by detection capability. Data processing occurs within Proofpoint's secured cloud infrastructure.

Customer Data and Model Training

Customer data is used in model training only where necessary and with appropriate controls.

Data Retention & Deletion

- Retention policies are automatically enforced by scheduled processes
- Data deletion is executed systematically in accordance with defined retention schedules

Data Access Controls

Access to data within AI pipelines is governed by role-based access controls with strict separation between pipeline, engineering, and data science roles. All privileged access is audited and logged.

Encryption

Data is encrypted at rest using dedicated encryption keys rotated on a defined schedule.

Customer Data Segregation

All customer data is stored with reference to customer tenant ID, enabling straightforward data management, deletion, and future segregation capability.

---

13. Third-Party & Cloud AI Services

Proofpoint's AI systems operate within managed cloud environments (AWS, Azure, GCP). Where Proofpoint utilizes cloud-hosted AI/ML services, these operate within Proofpoint's managed cloud tenancy under the following controls:

- Data processed through cloud-hosted AI/ML services remains within Proofpoint's managed cloud environment
- Cloud provider agreements prohibit use of customer data for provider model training
- Cloud-hosted AI/ML services are governed by the same security controls, access management, and data handling terms as all other cloud infrastructure
- Proofpoint discloses its use of cloud infrastructure providers and provides assurances about data security within those environments

Where open-source model architectures are utilized, model weights are hosted and managed within Proofpoint's own infrastructure.

---

14. Human Oversight

- Model promotion decisions require human review and approval across all detection systems
- Analyst teams continuously review model outputs, generating false positive and false negative signals that feed directly into model evaluation and retraining
- Human labelers are involved in training data quality processes where applicable, with access scoped to labeling tasks only

---

15. Versioning & Lineage

- Model versioning is maintained across all detection systems
- Rollback to previous model versions is supported
- Internal model documentation — including architecture specifications, deployed model metadata, and versioning history — is maintained in Proofpoint's internal systems
- Formal data lineage documentation is in development as part of Proofpoint's ISO 42001 certification.

---

16. Continuous Improvement

Proofpoint is committed to continuous improvement of its AI governance practices. Current investment areas include automated data quality assessment, automated pre-deployment validation, data lineage and provenance tracking.

This document reflects Proofpoint's AI data governance practices as of June 2026. Proofpoint is committed to continuous improvement and will update this document as practices evolve.