STANDARD SECURITY MEASURES

Last Updated: 2026-03-09

Version 1.0

This document describes Blueshift's standard security measures for informational purposes only and does not amend, expand, or otherwise modify any agreement between Blueshift and Customer, including the PLATFORM MASTER SUBSCRIPTION AGREEMENT, the DATA PROCESSING ADDENDUM — INTELLITHREAT AI PLATFORM, any applicable Order, or other applicable Addenda. Unless a specific security commitment in this document is expressly incorporated by reference into an applicable governing agreement, this document is non-binding and provided solely as a general description of Blueshift's then-current security practices. In the event of any conflict between this document and an applicable governing agreement, the applicable governing agreement controls.

1. Security Governance. Blueshift maintains commercially reasonable administrative, technical, and physical security controls designed to protect Customer Data and Personal Data against unauthorized or unlawful access, disclosure, alteration, or destruction, consistent with industry standards for cybersecurity software providers. Blueshift also maintains a written information security program and performs periodic, risk-based assessments and reviews of its security program.

2. Access Control. Multi-factor Authentication: Blueshift implements multi-factor authentication or other commercially reasonable access controls for personnel access to Blueshift systems used to provide the IntelliThreat AI Platform and to Customer Data, as appropriate to the nature, sensitivity, and risk profile of the access and system.

- **Least Privilege:** Access is granted based on the principle of least privilege and is reviewed periodically.

3. Data Protection. Encryption at Rest: Customer Data is encrypted at rest using industry-standard encryption measures, as appropriate to the relevant systems and data.

- **Encryption in Transit: Customer** Data is encrypted in transit using industry-standard encryption protocols, as appropriate to the relevant connection.
- **Separation:** Blueshift maintains measures designed to logically separate Customer Data and help prevent unauthorized cross-tenant access.

4. Vulnerability Management. Scanning: Blueshift performs periodic, risk-based automated vulnerability scanning of relevant in-scope production systems.

- **Penetration Testing:** Blueshift performs periodic internal and external security assessments, which may include penetration testing, vulnerability validation, and related technical reviews, based on risk, system scope, and operational considerations. Blueshift may make additional information regarding assessment practices available through its Trust Center or upon request, subject to confidentiality and security restrictions.

- **Patching:** Security patches and other remediation measures are applied based on risk, severity, availability, and operational considerations. Known Exploited Vulnerabilities (KEVs), as designated by CISA or an equivalent authoritative source, are targeted for remediation within thirty (30) days after an applicable patch or other commercially reasonable remediation becomes available, subject to prioritization, validation, testing, deployment windows, and operational constraints.

5. Physical Security. Blueshift maintains commercially reasonable physical security measures for facilities used in connection with providing the IntelliThreat AI Platform, which may include facility access controls, monitoring, and visitor management, as appropriate.

6. Logging and Monitoring. Blueshift maintains logging and monitoring controls for relevant production systems, including the IntelliThreat AI Platform and internal administrative systems. Unless otherwise expressly stated in an applicable governing agreement, logs are generally retained for a minimum of ninety (90) days in active storage and for a minimum of one (1) year in archival storage, in each case subject to applicable legal, security, system-design, and operational requirements. Blueshift employs SIEM-based monitoring to detect anomalous access, privilege escalation, and other security-relevant events on internal systems. Privileged administrative access to production environments is logged and subject to periodic review.

7. Incident Response. Blueshift maintains a documented incident response plan covering preparation, detection, containment, eradication, recovery, and post-incident review. Security incidents are internally classified by severity, and Blueshift maintains response objectives aligned to the severity and circumstances of the incident. The incident response plan is reviewed and tested at least annually. Customer notification obligations in connection with Personal Data incidents are governed by the DATA PROCESSING ADDENDUM — INTELLITHREAT AI PLATFORM §10 and are separate from Blueshift's internal incident response process described here.

8. Employee Security. Blueshift conducts background checks, where appropriate to the relevant role and level of access, on employees and contractors with access to production systems or Customer Data, to the extent permitted by applicable law. All personnel with such access are required to complete security awareness training at onboarding and at least annually thereafter. Access to Customer Data and production systems is provisioned based on role and reviewed upon role change, leave of absence, and termination. All Blueshift personnel are subject to confidentiality obligations no less protective than those in the applicable governing agreement.

9. AI Model Security. Given that the IntelliThreat AI Platform uses machine learning and agentic automation to support threat detection and response, and that customers who enable an Autonomous Mode tier may permit authorized Autonomous Actions within Customer-configured permissions, Blueshift implements the following additional security controls specific to AI model integrity, safety, and governance:

- **Model Integrity:** AI models used in the Platform are subject to internal validation and review appropriate to the nature, scope, and risk of the model or update before

deployment to production. Model updates undergo internal testing and, where appropriate, security and safety review prior to release.

- **Adversarial Input Testing:** Blueshift performs periodic, risk-based testing designed to evaluate Platform resilience against adversarial inputs, prompt injection attempts, and model manipulation techniques, consistent with commercially reasonable practices for AI-based security software.
- **AI Output Logging:** AI-generated detections and recommendations, and — for customers with an Autonomous Mode tier enabled — Autonomous Actions, are logged. Unless otherwise expressly stated in an applicable governing agreement, Autonomous Action logs are maintained as tamper-evident Audit Log records in accordance with Blueshift's then-current logging and retention practices, including any applicable requirements of the INTELLITHREAT AUTONOMOUS ADDENDUM where incorporated.
- **No Training on Customer Data:** Blueshift does not use identifiable Customer Data or Personal Data to train public, shared foundation, or third-party models and does not use identifiable Customer Data or Personal Data to train or fine-tune Blueshift's internal AI models except as expressly permitted by the applicable governing agreement or with Customer's prior written consent. Blueshift may use de-identified, anonymized, or aggregated telemetry data for model improvement to the extent permitted by the applicable governing agreement and applicable law.
- **Model Update Governance:** Material updates to core detection or action-execution models are subject to internal validation, testing, and security review, and, where appropriate, safety review before production deployment. Blueshift maintains commercially reasonable version control and rollback capabilities for production AI models, as appropriate to the relevant model architecture and deployment environment.

10. Autonomous Action Security Controls. For customers with the INTELLITHREAT AUTONOMOUS ADDENDUM incorporated and an applicable Autonomous Mode tier activated on an applicable Order, Blueshift implements the following additional security controls governing the execution of Autonomous Actions within Customer-configured integration environments:

- **Authorization Validation:** Before executing any Autonomous Action, the Platform is designed to validate, using commercially reasonable controls and the then-current capabilities of the applicable offering, that the action falls within the scope of Customer-configured permissions, scope definitions, exclusion lists, asset class permissions, and environment boundaries, and does not exceed the boundaries set in Customer's applicable configuration. Actions determined by the Platform to be outside configured scope are intended to be blocked and logged.
- **Rate Limiting:** Autonomous Action execution is subject to configurable or system-defined rate limits designed to help prevent runaway automation. Any customer-specific rate limit thresholds are documented in the applicable Order or in applicable Documentation only to the extent expressly incorporated by reference in the applicable governing agreement.
- **Cross-Tenant Isolation: Blueshift maintains commercially reasonable measures designed to logically isolate** Autonomous Action queues and execution

contexts on a per-Customer basis so that actions authorized by one Customer are not intended to affect the environment of any other Customer.

- **Kill-Switch Architecture:** Blueshift maintains a Kill-Switch capability that, when activated, is designed to disable or suspend Autonomous Mode as promptly as reasonably practicable and generally targets cessation of future or pending Autonomous Mode actions within sixty (60) seconds, subject to then-current system availability, network conditions, and platform state. Kill-Switch activation is logged. Kill-Switch activation for safety or security reasons constitutes Permitted Downtime under the SERVICE STATUS PAGE AND AVAILABILITY SLA only to the extent that SLA is applicable under the governing agreement.
- **Audit Log Replication: Blueshift is designed to replicate** Autonomous Action audit logs to a geographically separate secondary region on a reasonably prompt basis, and generally targets replication within twenty-four (24) hours of generation, subject to then-current system design, availability, and operational conditions, to support integrity and availability of audit records.