TRUST CENTER – SHORT FORM

Last Updated: 2026-03-09
Version 1.0

Security. Privacy. Transparency.

Blueshift is committed to protecting the confidentiality, integrity, and availability of Customer Data across the **IntelliThreat AI Platform** and supporting systems. This Trust Center – Short Form is provided for informational purposes only and summarizes certain security practices, privacy disclosures, compliance efforts, AI and autonomous feature disclosures, and operational safeguards. It does not create any contractual commitment, representation, warranty, service level, or other binding obligation unless and solely to the extent expressly incorporated by reference into the applicable governing agreement, and in the event of any conflict between this Trust Center – Short Form and an applicable governing agreement, the governing agreement controls. For clarity, this Trust Center – Short Form is directed primarily to the IntelliThreat AI Platform and related purchased offerings and does not by itself describe or create commitments for Blueshift's separate Managed SOC, Managed XDR, SOC-monitored, or other managed services except to the extent expressly stated in an applicable governing agreement.

---

1. Security Program Overview

Blueshift maintains a security program designed with reference to industry standards and practices, including the following:

- **NIST CSF**
- **SOC 2 readiness activities (in progress)**
- **Secure SDLC, internal review, and testing**
- **Continuous monitoring and vulnerability management**
- **Leastprivilege access controls**
- **Encryption in transit using TLS 1.2+ or successor industry-standard protocols, and encryption at rest using industry-standard measures appropriate to the relevant systems and data (including AES-256 where applicable)**

---

2. Data Protection & Privacy

Blueshift protects Customer Data and Personal Data processed in connection with the IntelliThreat AI Platform through measures such as the following, as applicable to the purchased offering, the relevant processing activity, and Blueshift's role under the governing agreement. For clarity, Personal Data processing associated with Blueshift's separate Managed SOC,

Managed XDR, SOC-monitored, or other managed services may be subject to separate agreements, disclosures, or addenda unless expressly made subject to the same governing terms.

- Encryption at rest and in transit
- Measures designed to support logical segregation in multitenant environments, as applicable
- Strict access controls and audit logging
- Data minimization and retention practices consistent with applicable agreements, legal requirements, and operational needs
- Subprocessor oversight and contractual safeguards
- Blueshift does not sell or share Personal Data, and does not retain, use, or disclose Personal Data except as permitted by the applicable governing agreements and law, including as necessary to provide, operate, maintain, secure, support, and improve the IntelliThreat AI Platform in accordance with such agreements and law.
- No use of identifiable Customer Data or Personal Data as such to train, fine-tune, or improve AI models; Blueshift may use de-identified, anonymized, and/or aggregated data only as permitted by applicable agreements and law, and does not use Personal Data as such to train public or third-party models

See the **DATA PROCESSING ADDENDUM – INTELLITHREAT AI PLATFORM (“DPA**”), if and where applicable, for additional details regarding Personal Data processing, international transfers, subprocessors, autonomous and AI-assisted processing disclosures, data use limitations, breach notification, and retention. HIPAA-related obligations apply only where expressly agreed in a signed contract, and Blueshift is not a default business associate or default BAA provider. Managed SOC, Managed XDR, and other managed-service privacy terms may be addressed in separate governing documents where applicable.

---

3. AI Safety & Responsible Use

The IntelliThreat AI Platform uses machine learning and agentic automation to support detection and response. Where Customer has activated or enabled an applicable IntelliThreat Autonomous offering, the Platform may execute authorized response actions autonomously within Customer-configured permissions, subject to the applicable configuration, consent, activation, and control framework.

- No training, fine-tuning, or improvement of AI models using identifiable Customer Data or Personal Data as such; Blueshift may use de-identified, anonymized, and/or aggregated data only as permitted by applicable agreements and law, and does not use identifiable Customer Data or Personal Data to train public or third-party models
- AI outputs and, where applicable, executed actions are logged and auditable, and Autonomous Actions include tamper-evident audit logs to the extent described in the applicable addendum

- Customer retains control over whether and how automation levels are enabled, configured, and constrained within the applicable purchased offering
- Configurable human oversight workflows, where available and applicable within the purchased offering
- Guardrails and controls designed to help prevent unauthorized actions
- Where the applicable autonomous offering is activated, Autonomous Actions are constrained by Customer-configured permissions, with Kill-Switch capability and tamper-evident audit logs (see the INTELLITHREAT AUTONOMOUS ADDENDUM)

---

4. Infrastructure Security

The IntelliThreat AI Platform is hosted primarily on **Oracle Cloud Infrastructure (OCI)** in U.S. regions and supported by infrastructure security measures such as:

- Physical security controls
- Redundant networking and power
- Continuous monitoring
- Hardened OS images
- Periodic internal and external security assessments, which may include penetration testing and related reviews, as appropriate in support of Blueshift's security program

---

5. Incident Response

Blueshift maintains a documented incident response plan that includes:

- 24/7 monitoring of Blueshiftowned systems
- Rapid triage and containment procedures
- Customer notification, as required by applicable law or the applicable governing agreement, following confirmation of a Security Incident affecting Personal Data without undue delay where and as required under such law or agreement; notice obligations are governed by the applicable agreement and are not triggered solely by unconfirmed suspicion of a potential incident
- Rootcause analysis and corrective actions

---

6. Compliance

Blueshift maintains controls designed to support compliance efforts, customer requirements, and risk management practices relating to the following, as applicable to the purchased offering, the relevant processing, and Blueshift's role under the governing agreement:

- SOC 2 readiness activities (in progress)
- GDPR/EEA data protection requirements, where applicable to Blueshift's role in the relevant processing
- CCPA/CPRA, where applicable to Blueshift's role in the relevant processing
- HIPAA-related obligations apply only where expressly agreed in a signed contract; Blueshift is not a default business associate and does not provide a BAA unless separately negotiated and executed
- NIST CSF (formal control mapping in progress; target Q4 2026)
- NIST AI RMF alignment considerations may inform certain AI governance and safety practices, but Blueshift does not represent that the IntelliThreat AI Platform is certified, validated, or formally compliant under any particular AI risk management framework unless expressly stated in applicable written documentation; AI governance materials made available in the Trust Center are generally informational unless expressly incorporated into a governing agreement

---

7. Subprocessors

The then-current subprocessor list for the IntelliThreat AI Platform is available in the **SUBPROCESSOR LIST – INTELLITHREAT AI PLATFORM at https://blueshiftcyber.com/legal/subprocessors/intellithreat-ai-platform. That list is provided for informational purposes and has only the notice, objection, and other limited procedural effect expressly stated in the applicable governing agreement documents; it does not by itself amend or expand any agreement**. For clarity, that list applies to the IntelliThreat AI Platform and does not identify subprocessors used exclusively for Blueshift's Managed SOC, Managed XDR, or other managed services, which remain subject to separate agreements and disclosures.

---

8. Responsible Disclosure

Security researchers may report vulnerabilities through the Responsible Disclosure & Vulnerability Reporting Policy. Reports should be submitted to security@blueshiftcyber.com. Blueshift generally seeks to acknowledge qualifying reports within a reasonable period and may apply coordinated disclosure timelines based on severity, exploitability, remediation complexity, and related risk considerations. Blueshift does not currently operate a monetary bug bounty program.

9. Standard Security Measures

Blueshift's detailed technical and organizational security controls — including AI model security, Autonomous Action security, logging, retention-related security control references, incident response, and employee security — are described in the STANDARD SECURITY

MEASURES document, as made available through Blueshift's public-facing legal or Trust Center materials, to the extent expressly incorporated by reference into applicable agreements. Unless expressly stated otherwise in the applicable governing agreement, those materials are informational only, describe the IntelliThreat AI Platform, and do not by themselves establish separate commitments for Managed SOC, Managed XDR, or other managed services.