Versa Responsible AI Disclosure

Last Updated: Jun 10 2026

Our Commitment

Versa Networks is committed to the responsible, secure, and transparent development and use of Artificial Intelligence and Machine Learning (AI/ML) across our unified SASE platform. We align our AI governance program with three leading frameworks:

- EU AI Act (Regulation 2024/1689)
- NIST AI Risk Management Framework (AI RMF 1.0)
- ISO/IEC 42001 — Artificial Intelligence Management System

This disclosure summarizes the AI systems embedded in Versa products, how customer data is handled, and how we meet transparency and governance obligations. For detailed technical information about any specific AI system, please contact ***aigovernance@versa-networks.com****.*

AI Ethics Statement

Our commitment to responsible AI is built on six core principles:

- **Fairness & Inclusion** — We actively work to identify and reduce bias in AI systems, ensuring decisions are equitable and do not discriminate.
- **Transparency** — We design AI systems to be understandable, with clear explanations for how outcomes are generated.
- **Privacy & Data Usage** — We apply Privacy-by-Design principles and use customer data only for agreed, legitimate purposes — never for unauthorized model training, resale, or unrelated product development. All data handling complies with GDPR, CCPA, the EU AI Act, and other applicable laws.
- **Security-by-Design** — We protect AI systems against misuse, adversarial attacks, and unauthorized access throughout their lifecycle — from data ingestion to deployment.
- **Accountability** — Every AI solution we deploy has clear human oversight, documented governance, and continuous monitoring to ensure responsible use.

- **Continuous Improvement** — We regularly review and update our AI practices in response to emerging risks, regulatory changes, and evolving best practices.

Our Pledge:

We will only develop, deploy, or partner with AI solutions that align with these principles and applicable laws. We regularly review our AI practices to adapt to evolving standards, emerging risks, and stakeholder expectations.

1. AI Systems in Versa Products

All Versa AI systems are embedded features within the VersaONE and Concerto platform. Based on the intended purpose and use-cases, none are classified as Unacceptable or High Risk under the EU AI Act.

Product	Purpose	Deployment	EU AI Act Classification
Verbo	AI assistant for troubleshooting, configuration verification, and natural-language commands in Director/Concerto. Includes Zero Trust MCP server.	Versa-hosted or On-premises. Requires Versa head-ends+MCP server	Minimal Risk
VANI	Predictive network analytics, anomaly detection, and intelligent alarm management	Versa-hosted or On-premises. Requires Advanced Security Cloud, Versa Messaging Service	Minimal Risk
UEBA	Behavioral anomaly detection including impossible travel, bulk operations, and MITRE ATT&CK patterns	Versa-hosted or On-premises. Requires Advanced Security Cloud, Versa Messaging Service	Limited Risk (transparency obligations)
ATP	AI-powered malware and threat detection across executables like PE, ELF, Javascript, VBScript, Office files like docx, pptx, xlsx, PDF and other file types	Versa-hosted or On-premises. Requires Advanced Security Cloud	Minimal Risk
AI DLP	AI-enhanced detection for source code, PII, images, document fingerprints, and OCR	Versa-hosted or on-premises through Versa Operating System VOS in SASE gateways	Limited Risk (transparency obligations)

Note: *EU AI Act classifications reflect cybersecurity and network security deployment contexts. Based on Versa's current classification of its intended use cases, Versa AI systems are not subject to the conformity assessment and Fundamental Rights Impact Assessment (FRIA) requirements applicable to High-Risk AI systems under Articles 27 and 43 of the EU AI Act . Classifications may vary based on customer-specific use cases: UEBA deployments used for employee monitoring, insider threat detection, or workforce management decisions may qualify as High Risk under Annex III and require conformity assessment, human oversight measures, and technical documentation. AI*

DLP deployments processing biometric-adjacent data (passport images, ID documents) should be reviewed independently.

Customers are advised to conduct and rely upon their own EU AI Act and data protection compliance assessment (including Data Protection Impact Assessments, where applicable under GDPR or other applicable laws) based on their intended use case and organizational context.

2. AI, Data Handling, Security and Tenant Isolation

- **Type of AI:** All Versa AI products use traditional machine learning techniques (also called as predictive AI). Verbo alone uses generative AI LLMs for response generation.
- **Multi-tenant isolation by design.** Every Versa AI system operates with full tenant separation across control, data, and management planes. Models in VANI, UEBA, and (where customer-specific) Verbo are trained and scoped per tenant with no cross-tenant data sharing.
- **Customer data is not used for generic model training.** ATP and AI DLP generic models are trained on industry datasets and synthetic data.
- **PII handling.** AI DLP scans PII only at inference for detection purposes; PII is never stored or used for training. UEBA processes identifiers such as names, emails, and IP addresses to correlate user activity, which is why it is classified as Limited Risk.
- **Encryption and access control.** All data is encrypted at rest and in transit. Cloud access is role-based with full audit trails.
- **Cloud inference:** Versa-hosted for hosted deployments; on-premises deployment for sovereign SASE customers.
- **Who has access to the data in Versa:** Only authorized Versa employees have access to this data on a need-to-know basis based on the principle of least privilege. **Data can only be viewed and cannot be exfiltrated out of the AI systems.**

3. Explainability, Testing, and Audit

- Dedicated explainability models exist in Verbo, VANI, and UEBA.
- AI DLP uses task-specific models producing explainable outputs per detection type.
- ML lifecycle management tool tracks all model experiments, versions, metrics, and transactions, providing a complete audit trail.
- Chat feedback signals (Verbo) and model drift detection based on efficacy percentages, feedback attribution (ATP) provide ongoing quality assurance.

4. Framework Alignment

Framework	Versa Alignment
EU AI Act (Regulation 2024/1689)	Risk classification applied per AI system. All systems classified as Minimal or Limited Risk. Transparency obligations met for UEBA and AI DLP. No High Risk systems — Conformity Assessment and Fundamental Rights Impact Assessment (FRIA) not required.
NIST AI RMF 1.0	Operationalized across the Govern, Map, Measure, and Manage functions. Model inventory, risk assessment, explainability, and continuous monitoring via ML lifecycle management tools.
ISO/IEC 42001	AI Management System alignment across governance, lifecycle controls, and third-party model oversight. Certification roadmap in progress.

5. Third-Party Models and Sub-Processors

Versa uses a combination of proprietary and third-party models. Model versions may be swapped or upgraded without notice; the current inventory is maintained in the Versa Trust Center and can be obtained from aigovernance@versa-networks.com based on request

6. Governance and Oversight

Versa's AI governance program is overseen by Product, Engineering, Privacy, and Legal stakeholders and is reviewed periodically. All AI decisions (such as block and quarantine actions and UEBA risk scoring) include administrator review, override, and audit logging.

7. Request Additional Information

For detailed technical information about any specific AI system — including model architecture, training data sources, evaluation metrics, deployment specifics, or per-product transparency documentation — please contact our AI governance team:

aigovernance@versa-networks.com

Customers and researchers may also raise security concerns via the Versa PSIRT process.