Privacy Statement

Version 1.0 · 25 June 2026

1. About this privacy statement

This privacy statement describes how Attic Security B.V. ("Attic", "we", "us") processes personal data in the context of our MDR services and our relationship with our customers and users.

We value your privacy and handle personal data in accordance with the General Data Protection Regulation (GDPR) and the Dutch GDPR Implementation Act (Uitvoeringswet AVG).

This statement applies to:

- Customer organisations using our services, including continuous monitoring (DSCM), IVON (our AI-driven response agent) and related MDR components
- End users within customer organisations whose data is processed
- Visitors to our website and office
- Individuals who contact us for sales, support or other purposes

---

2. Who we are

Attic Security B.V. Molenstraat 36, 4761 CL Zevenbergen, the Netherlands Chamber of Commerce (KvK): 83973206 Email: privacy@atticsecurity.com Website: www.atticsecurity.com

Privacy contact: Erik Remmelzwaal, CEO privacy@atticsecurity.com

---

3. What data we process

We process two categories of personal data, each with its own processing context.

3.1 Service data — our MDR services

To deliver our MDR services we process personal data from our customers' Microsoft 365 environment (tenant). Our services comprise several components:

- **Continuous monitoring and posture management** — periodic security checks on your Microsoft 365 environment to detect anomalies, misconfigurations and risks
- **IVON** — AI-driven triage and remediation proposals for security incidents (see section 5 for AI-specific detail)
- Additional MDR components as agreed in your Agreement

As the customer, you remain the controller of this data at all times; Attic acts as processor, in accordance with the data processing agreement (DPA) concluded between us.

Specific categories of data we process on the customer's instructions:

- **Identity and authentication data** from Microsoft Entra ID — user names, email addresses, role information, sign-in history, risk events
- **Device and session information** from Microsoft Defender — endpoint IDs, IP addresses, session tokens, threat signals
- **Security incidents and alerts** from Microsoft Sentinel and Defender — investigation data, alert details, indicators of compromise
- **Audit log data** from Microsoft 365 — user actions relevant to security investigations
- **Configuration and posture data** — security settings, compliance status, access configuration (for monitoring components)

We access this data exclusively through delegated OAuth tokens that you, as the customer, provide to us via Microsoft Entra ID. Access is isolated per customer tenant; our systems technically reject any attempt at cross-tenant access.

You can revoke this access at any time via Microsoft Entra ID.

3.2 Customer relationship data

For our customer relationship, sales, support and billing, we process personal data of contacts at customer organisations and prospects:

- **Contact details** — name, job title, business email address, telephone number, organisation
- **Communication history** — emails, support tickets, conversation notes
- **Billing data** — company name, billing address, payment information (processed via Chargebee and Stripe)
- **Marketing interactions** — where applicable and subject to your preferences

For this category we are the controller.

---

4. Why we process data

Category	Purpose	Legal basis (GDPR)
Service data (3.1)	Delivering the MDR service under the agreement and data processing agreement	Art. 6(1)(b) (performance of a contract, in the processor role)
Customer relationship data (3.2)	Managing the business relationship, service delivery, support, billing	Art. 6(1)(b) (contract) and Art. 6(1)(f) (legitimate interest)
Marketing communication	Informing you about services and products	Art. 6(1)(a) (consent) or Art. 6(1)(f) (legitimate interest, opt-out available)
Legal obligations	Statutory accounting and tax retention requirements	Art. 6(1)(c) (legal obligation)

---

5. How we process data — IVON in detail

IVON is an automated analysis platform that identifies, triages and proposes remediation actions for security incidents in your Microsoft 365 environment. Because of the AI component, we explain here in detail how it works.

5.1 Four safeguards for responsible AI processing

Personal data of your end users is passed to our language model for incident analysis. This data is protected by four safeguards that we have established contractually and technically:

1. **Processing exclusively within the EU.** All AI inference takes place within Microsoft Azure AI Foundry in the Sweden Central region. Our own backend runs within Hetzner data centres in Germany and Finland. No transfer outside the European Economic Area.
2. **The language model is not trained on your data.** Microsoft Azure contractually guarantees that your prompts and responses are not used to train or fine-tune any AI model.
3. **Prompts and responses are runtime-only.** The language model processes your data only during the active incident analysis. There is no persistent storage at the language model, no memory between sessions and no profiling of your organisation or end users.
4. **Isolated per customer, with no cross-contamination.** Access is via OAuth tokens specific to your tenant; our systems technically reject any attempt to retrieve data from other tenants. For periodic security checks we additionally apply strict runtime separation, where each check runs in a clean environment with no residue from earlier customer sessions, using centrally managed access credentials.

5.2 Data flow at a high level

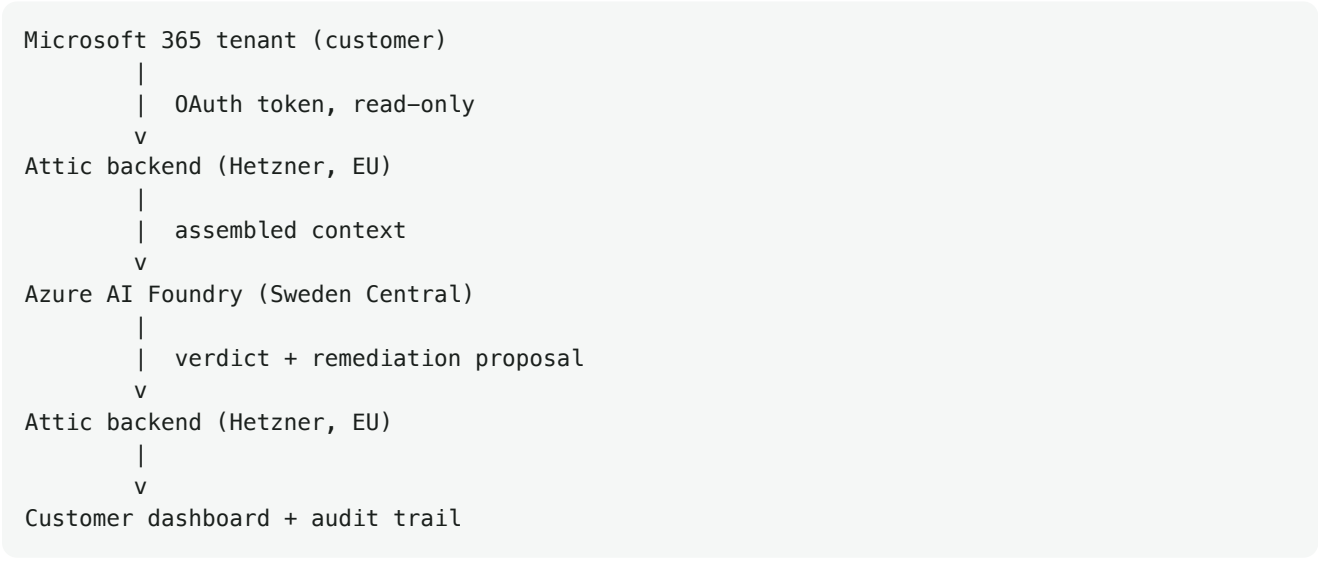

5.3 Processing steps

1. **Ingest.** We read incidents, alerts and relevant context from your Microsoft 365 environment via Microsoft Graph and related APIs. We do not modify your data.
2. **Prepare.** On our backend (Hetzner, EU) we assemble a relevant context for the specific incident analysis. This context package contains the alert data, the user names involved, IP addresses,

indicators of compromise and the specific question put to the language model.

3. **Analyse.** The assembled context is sent to the language model. The model produces a verdict (true positive, false positive, benign or inconclusive) and any remediation proposals. The language model sees only this single incident context — no history, no other customers, no access to the wider tenant.
4. **Return.** The result comes back to our backend and is made available to you in the Attic dashboard, including an audit trail.
5. **Execute** (optional). Remediation actions approved by you are carried out in your own tenant via Microsoft Graph. No changes are made to your environment without your approval.

5.4 Technical transparency — which language model

The language model we use is Claude, developed by Anthropic and hosted by Microsoft within Azure AI Foundry in the Sweden Central region. In this setup Microsoft is responsible for the hosting, data handling and associated contractual safeguards — Anthropic itself receives no customer data via this route. For that reason Anthropic is not a sub-processor of Attic; Microsoft Corporation (Azure AI Foundry) is, and is listed in our sub-processor list (Annex A).

5.5 AI processing — additional characteristics

In addition to the four core safeguards in section 5.1, the following technical characteristics are relevant:

- **Abuse monitoring by Microsoft.** Microsoft applies content filtering to AI input and output. Content flagged by classifiers may be retained in-region (Sweden Central) for up to 30 days for human review by Microsoft, logically separated per resource, and is not used for model training. This measure is intended to counter abuse of AI services — such as the generation of malware or harmful content.
- **Transient caching.** For performance, Azure applies transient prompt caching with a maximum of 24 hours, GPU-local and separated per Azure subscription.
- **No stateful AI functions.** We do not use Azure functions that build persistent AI state, such as Assistants threads, fine-tuning or vector indexes.

5.6 Audit trail

We retain an audit trail of every analysis performed on our own backend (Hetzner, EU). This audit trail contains metadata such as the time the alarm was received, the time to pick-up, enrichment steps, actions taken by IVON, the conclusion and the advice given to you. This audit trail is visible only to you and to Attic staff who carry out investigation work on your instructions.

5.7 Future privacy options and customisation

We continue to work on additional privacy controls. One option we are considering for future releases is pseudonymisation prior to AI processing, whereby certain personal data (such as user names and email addresses) would be replaced with pseudonyms before being sent to the language model. This functionality is **not yet available at launch** — we are evaluating it alongside customer feedback for introduction in a later release.

Do you have specific heightened privacy requirements arising from your sector or regulations (for example healthcare, government, legal services)? Contact privacy@atticsecurity.com to discuss the options for tailored arrangements.

6. Automated decision-making

IVON performs automated analysis, but no automated decision-making producing legal effects or similarly significant impact within the meaning of Art. 22 GDPR. Specifically:

- IVON triages and proposes actions; remediation actions are only carried out after explicit approval by you or a person authorised by you
- No automatic sanction or refusal is applied to data subjects on the basis of AI analysis

7. With whom we share data

To deliver our services we use carefully selected sub-processors, with each of whom a data processing agreement has been concluded. The full sub-processor list can be found in Annex A of this document and is updated periodically on our website.

We share your personal data only with:

- **Sub-processors** necessary for our service delivery (see Annex A)
- **Competent authorities** where a legal obligation to do so exists
- **Advisors and accountants** insofar as necessary for our business operations, under a confidentiality obligation

We do not sell personal data to third parties for marketing or other commercial purposes.

8. Where we process data — no transfer outside the EU

Our processing takes place entirely within the European Union:

- **Backend and storage:** Hetzner Online data centres in Germany (primary) and Finland (secondary)
- **AI inference:** Microsoft Azure AI Foundry in Sweden Central
- **Continuous monitoring infrastructure:** Microsoft Azure region West Europe
- **Customer relationship data (CRM, support, billing):** EU regions of our sub-processors — see Annex A

We do not transfer personal data to countries outside the European Economic Area (EEA). Should such a transfer become necessary in the future, it will take place exclusively under the applicable safeguards of Chapter V GDPR (such as Standard Contractual Clauses) and this privacy statement will be amended beforehand.

9. How long we retain data

Category	Retention period	Notes
Service data — IVON incident investigation data	**30 days** in IVON system	Enables us to investigate quickly at your request. Passed to continuous monitoring systems for longer retention where relevant
Service data — continuous monitoring (data selected by checks)	**1 year**	For historical overview and posture trend analysis
Service data — open alarms and outstanding tickets	Until resolution + 1 year	Until the workflow is complete
Service data — on termination of a customer subscription	7-day soft delete + 30 days thereafter purge	Full deletion within 37 days of subscription termination
Service data — LLM prompts/responses	Runtime-only	See section 5.5
Service data — content flagged by Microsoft abuse monitoring	Maximum of 30 days at Microsoft	See section 5.5
Customer relationship data — billing-related	7 years	Statutory tax retention requirement
Customer relationship data — other (CRM, contacts)	Duration of the customer relationship + 24 months	For relationship management
Support tickets	Duration of the customer relationship + 24 months	For ongoing issues
Marketing data	Until consent is withdrawn or opt-out	See section 11
Website visitor information	In accordance with our cookie policy	atticsecurity.com/cookies

The original data that Attic works through (the incidents, alerts and log data in your Microsoft 365 environment) always remains your property and stays in your environment. After the periods above, Attic deletes only its own copies and derived processing.

---

10. How we secure data

Attic is **in the process of obtaining ISO/IEC 27001:2023 certification, with our Stage 2 audit scheduled for Q4 2026.** We apply appropriate technical and organisational measures to protect personal data, including:

- **Access security** — least-privilege access, multi-factor authentication for all staff with access to customer-relevant systems

- **Encryption** — encryption in transit (TLS) and at rest for all personal data
- **Logging and monitoring** — all access to customer data is logged and monitored; anomalies are investigated
- **Incident management** — we operate a formal incident management process with reporting routes. We report personal data breaches affecting your personal data to you within 48 hours; you then, as controller, report to the Dutch Data Protection Authority within 72 hours in accordance with GDPR Art. 33
- **Awareness** — our staff receive periodic training on information security and privacy
- **Sub-processor review** — we assess our sub-processors on security and privacy at contracting and periodically thereafter

---

11. Your rights

Under the GDPR, you (or the data subjects in your organisation) have the following rights:

- **Access** to the personal data we process
- **Rectification** of inaccurate or incomplete data
- **Erasure** (the right to be forgotten), insofar as there is no legal basis for retention
- **Restriction** of processing in certain circumstances
- **Portability** of data
- **Objection** to processing based on legitimate interest or for direct marketing
- **Withdrawal of consent** where processing is based on consent

How to exercise your rights: send an email to privacy@atticsecurity.com. We respond within 30 days.

For service data (category 3.1) for which your organisation is the controller: data subject requests are handled in the first instance by your own organisation. We support you in doing so in accordance with our data processing agreement.

---

12. Right to complain

If you believe that we are not handling your personal data with due care, we would like to hear from you at privacy@atticsecurity.com.

You have the right to lodge a complaint with the Dutch Data Protection Authority (Autoriteit Persoonsgegevens): **Autoriteit Persoonsgegevens** Postbus 93374 2509 AJ The Hague, the Netherlands www.autoriteitpersoonsgegevens.nl

---

13. Changes

We update this privacy statement when our services, sub-processors or relevant laws and regulations give cause to do so. The current version is always available at www.atticsecurity.com/privacy.

In the event of material changes we inform you proactively, in any case by email to the registered contact at your organisation.

14. Governing law and disputes

This statement is governed by Dutch law.

15. Contact

For questions about this privacy statement or the processing of your data:

- **Privacy questions:** privacy@atticsecurity.com
- **Privacy contact:** privacy@atticsecurity.com (Erik Remmelzwaal)
- **General:** contact@atticsecurity.com

Annex A — Sub-processor overview

Sub-processor	Data category	Processing location	Purpose
Hetzner Online GmbH	Service data (all, audit trail), customer configuration	Germany + Finland (EU)	Hosting backend infrastructure
Microsoft Corporation (Azure + Azure AI Foundry)	Service data (continuous monitoring, LLM input/output runtime-only, vault storage)	West Europe + Sweden Central (EU)	Cloud infrastructure and AI inference
Zendesk Inc.	Customer relationship data (support tickets)	EU region	Customer support and ticket management (planned migration to HubSpot)
HubSpot Inc.	Customer relationship data (CRM)	EU region	CRM and relationship management
Chargebee Inc.	Customer relationship data (billing)	EU region	Billing and subscription management
Stripe Inc.	Customer relationship data (payment)	EU region	Payment processing (via Chargebee)
Functional Software Inc. (Sentry)	Service metadata (error tracking, may contain PII in stack traces)	EU region	Error detection and debugging (planned replacement by the self-hosted GlitchTip)

Technical infrastructure (no personal data, not a GDPR sub-processor but listed for transparency):

- **Zabbix** — system monitoring, infrastructure metrics only

- **GlitchTip** — error tracking, self-hosted solution on Attic infrastructure (Hetzner, EU)