Data Processing Agreement

Sunbeam Control — AI Workforce Operations Platform

Issued by: Mishki Ltd trading as Sunbeam Control
Version: 1.0
Date: March 2026
Contact: legal@sunbeamcontrol.com

This Data Processing Agreement ("DPA") describes how Mishki Ltd trading as Sunbeam Control ("Sunbeam", "we", "us") processes data in connection with the Sunbeam Control platform. It is provided to enterprise procurement teams as part of vendor due diligence.

1. Overview

Sunbeam Control is an AI Workforce Operations Platform. It finds every AI tool, agent, and API connection operating in a customer's environment and provides the information needed to manage, audit, and remediate that workforce.

The fundamental design principle of Sunbeam is local-first data processing. Scan findings are generated on the customer's own infrastructure and are not transmitted to Sunbeam or any third party in the standard product configuration.

2. Data Controller and Processor Roles

2.1 Customer as Controller

The organisation deploying Sunbeam Control ("Customer") is the data controller for all data processed during a scan. The Customer determines the purpose and scope of each scan, the machines on which it runs, and the use of the resulting findings.

2.2 Sunbeam as Processor

Mishki Ltd acts as a data processor only where the Customer uses the optional Central Server feature (v2.5.0+), which collects scan findings from multiple machines and aggregates them in a self-hosted dashboard. In this configuration, Sunbeam processes data on behalf of the Customer in accordance with the Customer's documented instructions.

2.3 No Sunbeam Cloud Processing

In the standard configuration (standalone scanner, CLI mode, or single-machine deployment), Sunbeam does not act as a data processor. All processing occurs on the Customer's own machine. Sunbeam does not receive, store, or have access to any scan findings.

3. What Data Sunbeam Processes

3.1 Data Processed Locally (Standard Configuration)

The following categories of data are processed locally on the scanned machine and are not transmitted to Sunbeam:

- Network topology: local IP addresses, open ports, HTTP response headers and body snippets
- Process information: running process names, listening ports, outbound connection destinations
- Filesystem artefacts: browser history database queries (AI domains only), shell history entries (AI command patterns only), application log file existence and timestamps
- Environment variables: variable names matching AI API key patterns; values are detected but masked in all output
- VS Code extension manifests: extension identifiers matched against known AI tool signatures
- Installed application names: matched against known AI tool signatures
- Cloud API responses: service names and regions returned by AWS, Azure, and GCP APIs when credentials are present
- OS user attribution: the username of the account that installed or owns each detected agent

3.2 Data Processed by Central Server (Optional — v2.5.0+)

If the Customer deploys the optional Central Server, the following data is transmitted from each enrolled machine to the Central Server over the Customer's own network:

- Machine hostname
- Scan timestamp
- Findings in the standard 12-field JSON schema (see Section 5)
- Per-machine API key (for authentication — not a personal credential)

The Central Server runs on the Customer's own infrastructure. Data is stored in a SQLite database on the Customer's server. No data is transmitted to Sunbeam.

3.3 Data Sunbeam Never Collects

The following categories of data are never collected, transmitted, or stored by Sunbeam in any product configuration:

- The content of environment variable values (keys are detected; values are not recorded)
- The content of browser history entries beyond the domain name
- The content of shell history beyond AI-related command patterns
- Credentials, passwords, tokens, or private keys
- Personal communications, documents, or files
- Telemetry, usage statistics, or diagnostic data from the Customer's environment

4. Data Residency

Scan findings reside exclusively on the infrastructure designated by the Customer:

- Standard configuration: findings reside on the scanned machine only, in memory during the scan and in the output file (PDF, CSV, or JSON) if the Customer saves them.
- Central Server configuration: findings reside on the Customer's Central Server, in the Customer's chosen deployment environment (on-premises, private cloud, or Customer-controlled cloud tenancy).

Sunbeam does not operate any cloud service that receives Customer data. There is no Sunbeam SaaS backend for scan findings in any current product version.

5. JSON Output Schema

Sunbeam produces findings in a stable, versioned 12-field JSON schema. The schema is defined below and will not change without a version increment. Customers integrating Sunbeam output into downstream systems (SIEM, CMDB, data warehouse) may rely on this schema from v2.2.0 onwards.

Field	Type	Description
version	string	Schema version — e.g. "2.4.0"
timestamp	ISO 8601	Scan start time in UTC
machine	string	Hostname of the scanned machine
summary.total	integer	Total findings across all surfaces
summary.high	integer	Count of HIGH risk findings
summary.medium	integer	Count of MEDIUM risk findings
summary.low	integer	Count of LOW risk findings
findings[].surface	string	workstation | network | egress | cloud
findings[].service	string	Human-readable agent name
findings[].risk	string	HIGH | MEDIUM | LOW
findings[].owner	string	OS username attributed to this agent
findings[].remediation	string	Exact removal instructions for this agent

6. Sub-processors

In the standard product configuration, Sunbeam does not use sub-processors. No third-party services receive Customer data.

If the Customer uses the optional SIEM or webhook integration, the Customer configures the destination endpoint. Sunbeam transmits findings to the Customer-designated endpoint on the Customer's instruction. The Customer is responsible for ensuring that destination endpoint complies with applicable data protection law.

7. Security Measures

- All scan findings are processed in memory and written only to Customer-designated output paths
- The Central Server uses per-machine API keys for authentication; keys are revocable per machine
- Signed and timestamped scan evidence (v2.6.0+) uses SHA-256 and HMAC to provide tamper-evident chain-of-custody for audit purposes
- The Sunbeam scanner binary is code-signed (macOS: Apple Developer ID; Windows: DigiCert EV — pending)
- No credentials or private key material is stored in any Sunbeam output

8. Retention

Sunbeam does not retain Customer data. Scan output files (PDF, CSV, JSON) are created at the Customer's direction and retained by the Customer according to the Customer's own retention policies.

The Central Server (v2.5.0+) stores findings in a SQLite database on the Customer's server. The Customer is responsible for the retention and deletion of that data.

9. Data Subject Rights

Scan findings may include the OS username attributed to an installed AI agent. This constitutes personal data under GDPR where the username identifies a natural person.

The Customer, as data controller, is responsible for handling data subject access requests, erasure requests, and other rights requests relating to findings data. Sunbeam, in its capacity as processor under the Central Server configuration, will assist the Customer in fulfilling such requests upon written instruction.

10. Governing Law

This DPA is governed by the laws of England and Wales. Mishki Ltd is a company registered in England and Wales.

For data protection enquiries: legal@sunbeamcontrol.com