Policy for data ethics

Just Software ApS · Approved · Data ethics

Version	v1.8	**Approved on**	5 November 2026
Approval authority	CISO or CTO	**Next review due**	5 November 2027
Last reviewed	5 November 2026	**Review interval**	365 days
Relevant frameworks	D-mærket		

Policy for data ethics

1. Goals

Just Software builds infrastructure and security tools that process data about our customers' users, devices and cloud environments. With this policy we aim to:

- consider the ethical consequences of our products before and while we build them, not only their legal compliance;
- keep the data we process minimal, purpose-bound and under the customer's control;
- use AI only where the customer chooses it, transparently and without automated decisions about people;
- be open with customers and users about what we process and why.

2. Scope

The policy covers Just Software's outward-facing activities, i.e. all products offered to customers:

Product	Nature	Personal data processed (in short)
EasyScep	Cloud SCEP certificate authority for Intune-managed devices	Certificate subject/SAN data (names, UPNs, e-mail), device identifiers, request logs
EasyRadius	Cloud RADIUS (802.1X, EAP-TLS) network access control	User names/UPNs, certificate data, MAC addresses, network device identifiers and IP addresses, authentication outcomes
DNS Tool for Azure	Monitoring of Azure DNS zones, certificates and domains	Mainly technical DNS data; user accounts of customer administrators
Praefic	Azure security posture, cost and change analysis	Azure resource configuration and findings; Azure Activity Log identities (who changed what, from where) unless switched off by the customer
Sched4Cloud	Scheduling of Azure resources for cost savings	Technical configuration; user accounts of customer administrators
CSV Bridge	Windows application for importing CSV files into Excel	None — all processing happens locally on the user's device

For all hosted products, Just Software acts as **data processor** on the customer's (data controller's) documented instructions, as set out in the Data Processing Agreement (DPA) and its product annexes. Our direct users are the customers' IT administrators; the data subjects are mostly the customers' employees and other end users whose devices or accounts are processed.

The policy applies to all employees and to everyone involved in product development, operations, support and sales.

3. Principles

1. **Customer control.** Customer data is processed only to deliver the service, on the customer's instructions. We do not sell it or use it for our own purposes.
2. **Data minimisation.** We collect only what the product needs to work. Example: EasyRadius supports only certificate-based authentication (EAP-TLS) and therefore stores no passwords.

3. **Limited retention.** Data is kept for defined periods and then deleted or anonymised. Personal identifiers in operational and security logs are irreversibly removed after 90 days; minimised event data may be kept for up to 12 months for security purposes.
4. **AI by choice.** AI features that process customer data are off by default unless they are a documented core function, are enabled by the customer, and the data sent and the provider used are documented publicly. AI never makes decisions about people's access or rights.
5. **No profiling of people.** Our products do not evaluate, rank or treat people differently based on personal characteristics such as age, gender, ethnicity or religion.
6. **Transparency.** We describe in plain language what each product processes, where, for how long, and which sub-processors are involved.

4. Implementation

- **Product design and changes:** new features and changes that affect personal data or AI are assessed against the principles in section 3, as part of the *Privacy and security by design* policy and the *Software development* policy. AI use-cases are recorded in the AI use-case register in Oopsie.
- **AI:** transparency and review requests are handled as described in section 5, and stakeholder involvement as described in section 6.
- **Data subject rights:** as processor, we assist the controller. Customer administrators can search records per user or device and export them (CSV/JSON) in the product portals. Erasure takes place through fixed retention periods and deletion at the end of the subscription as described in the DPA. Correction does not apply to log and event records, which are factual records of events.
- **Support data:** handled under the *Anonymisation of support data procedure*.
- **Raising concerns:** any employee who finds a feature, process or use of data ethically unclear or potentially harmful raises it with the CTO or CISO before it is released. Concerns are discussed openly and the outcome is recorded.
- **Publication:** the content of the approved data ethics table (section 7) is published on the Just Software Trust Center.

5. AI transparency and review requests

- **Decision support only.** AI proposes; a person decides. Use-cases and the AI Responsible are recorded in the AI use-case register in Oopsie.
- **Labelling.** Every AI output is marked as AI-generated where a person sees it (product UI, exports, reports).
- **Plain language.** Each use-case is explained by what it looks at, what it produces, who decides and what it does not do.
- **Where people are informed.** The Trust Center (use of AI, AI providers as sub-processors, per-product data-protection statements, text customers can reuse for their end users), the Privacy policy for our own processing, and the product documentation.
- **Review requests.** As published on the Trust Center, anyone affected by AI output or an AI-supported decision can email security@just-software.com with the subject "AI review". We act within 10 working days: a person who did not take the original decision reviews the case without relying on the AI output and replies in plain language. A review need not lead to a change. Requests and outcomes are recorded and are input to the next stakeholder consultation and AI risk assessment.
- **Accuracy.** Before a new AI use-case or a major change (section 6) is released, the published texts are checked against the configuration and the vendor register.

6. AI stakeholder involvement

- **When.** For every AI use-case in the register, stakeholders are consulted before go-live, during any pilot or staged rollout, within three months after go-live, on every major change, and at least once a year together with the AI risk assessment.
- **Who.** The AI Responsible leads each consultation with the CISO. Because the people who build the use-cases cannot foresee every perspective, each consultation includes at least one voice from outside development — for example customer administrators, feedback from support tickets, objections and data subject requests, or a non-developer user.
- **Major change** means any of: a change of model or inference provider; a prompt or configuration change that materially alters output; a new data source or data category; a change to a default that affects personal data; a new customer segment with a different legal position.

- **Decisions.** Each input is accepted (with an action), rejected (with a reason) or deferred (with a date). Outcomes that change risk feed into the AI risk assessment; outcomes that change scope feed into the AI use-case register.
- **Records.** Each consultation is recorded in Oopsie's AI consultation records and signed off by the CISO. Records are kept while the use-case is in the register and for three years after it leaves.
- The DS/PAS 2500 question sets may be used as checklists.

7. Data ethics table

Status scale: Not relevant · No considerations · Some considerations · Actions planned · Actions implemented · Continuous process established.

#	Question	Status	Considerations
1	Has the company considered solutions to improve its work with data ethics (e.g. working group, awareness, culture of openness)?	Actions implemented	As a small company, management (CISO and CTO) acts as the data ethics forum. Data ethics is part of product design reviews, this table is reviewed at least yearly, and employees are expected to raise concerns openly (section 4).
2	Has the company considered ongoing involvement of relevant stakeholders, e.g. to avoid unintended bias?	Actions implemented	Stakeholders for AI use cases are identified and involved as described in section 6. Customer feedback through support and onboarding is used to adjust product behaviour.
3	Has the company considered the consequences of influencing user behaviour (e.g. addiction, overuse, bullying)?	Not relevant	Our products are infrastructure and security tools for IT administrators. They contain no engagement mechanisms, personalised content, nudging or social features. Notifications are configured by the customer.
4	If the company influences user behaviour, has it considered how to make this more transparent and give users more control?	Not relevant	See question 3.
5	Has the company considered how users' rights are prioritised over commercial or institutional interests?	Actions implemented	We process customer data only as processor on the customer's instructions and do not sell it or use it for our own purposes. AI features that process customer data are off by default and enabled by the customer. Products are hosted in the EU unless the customer chooses otherwise.

#	Question	Status	Considerations
6	Has the company considered how users get the most value from the data collected?	Actions implemented	Data is used to give customers insight into their own environment: authentication and certificate logs with diagnostics, DNS change history and findings, posture and cost reports. Customers can search and export their data (CSV/JSON).
7	Has the company considered how to avoid unintended consequences (e.g. surveillance, misuse, spreading misinformation)?	Actions implemented	The main risk is that logs of who connected, when and from which device (EasyRadius, EasyScep), or who changed cloud resources (Praefic), could be used to monitor employees. Mitigations: data minimisation, limited retention with identifiers removed after 90 days, Praefic identity collection can be switched off by the customer, and access to logs limited to the customer's authorised administrators. AI output (e.g. DNS change risk scores, diagnostics, report narratives) is advisory only, never changes customer systems, and the data sent and the provider are documented. Where the customer chooses an AI provider outside the EU, this is stated in the product documentation.
8	Has the company considered how to protect special target groups (e.g. children, young people, groups with special challenges)?	Not relevant	Our products are business-to-business infrastructure tools used by organisations' IT administrators and are not directed at children or other vulnerable groups.
9	Has the company considered whether data collection and solutions/products can limit citizens' rights?	Actions implemented	EasyRadius makes network access decisions and EasyScep can revoke certificates, which can deny a person network access. These decisions follow rules configured by the customer only; no AI or automated profiling is involved. Failure reasons are shown to administrators so that wrongly denied users can be helped quickly.

#	Question	Status	Considerations
10	Has the company considered avoiding reinforcing social and ethical problems (e.g. inequality, exposing population groups, solutions usable only by certain users)?	Some considerations	Products contain no logic based on personal characteristics (principle 5). AI use cases are reviewed for bias in the stakeholder consultations (section 6). The products depend on the Microsoft ecosystem (Entra ID, Intune, Azure), which is inherent to their purpose.
11	Does the company communicate its privacy-by-design strategies to its users?	Actions implemented	Communicated through the public Trust Center (sub-processors, hosting, how AI is used), per-product data-protection statements, product documentation FAQs (e.g. "Data & AI"), and the DPA with product annexes.

8. Follow-up and reporting to management

- The data ethics table is reviewed and updated at least once a year, and whenever there are major changes (e.g. a new product, a new AI use case or provider, changed data flows).
- The updated table is approved by management and the public version on the Trust Center is updated.
- Ethical concerns raised, their outcome and the status of this table are reported in the *Management review*.

How is this policy implemented and validated?

Recurring tasks for security roles

Annual AI stakeholder consultation

- Description: Consult stakeholders on every AI use-case in the register, as described in section 6.
- Frequency: Yearly
- Security role: AI Responsible (CTO)

Data ethics table review

- Description: Review the data ethics table, have it approved by management and update the Trust Center.
- Frequency: Yearly
- Security role: CISO