

Specimen — illustrative example

Solstice Grid Services is a fictional organisation. This document comes from the scenario described on riskwoven.com/examples, and is not an assessment of any real entity. It is advisory output to support your own decisions: not legal advice, not a certification, and not an audit.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.

Response: Solstice supplier AI questionnaire

Solstice Grid Services · Written and reviewed by RiskWoven.com

Requested by: Ontario grid partner · **Questions:** 4 · **Answered:** 2 · **Partly answered:** 2 · **Not answered:** 0

Each answer below is drawn only from the documents and evidence in the workspace. 2 questions name the evidence that would complete it.

1. Do you have a documented AI policy?

Answered

We have a documented AI Governance and Accountability Policy. It establishes the accountability framework for all AI systems we use, mandating that every system is inventoried, risk-assessed, assigned a named owner, subject to proportionate human oversight, monitored throughout its operational life, and retired securely. The policy covers the full lifecycle from proposal through to retirement and applies to every employee, contractor and third party who proposes, develops, deploys, operates, monitors or retires an AI system on our behalf. The policy is currently documented but not yet in effect, as it is at version 1.0 Draft pending release.

Source: AI Governance and Accountability Policy, 1. Purpose and 2. Scope

2. How do you assess the impact of AI systems on customers?

Answered

Before any AI system is deployed, and before any material change to an existing system, we conduct a documented impact assessment. The assessment addresses the business purpose and intended use, the data the system will receive including any personal information, the outputs or actions the system will produce, the potential harms if the system fails, drifts, is misused or is compromised, the controls that reduce those harms, the residual risk after controls, and the human oversight arrangements. For systems that affect customers, the assessment specifically confirms that no customer data will be used for vendor model training and that data will not leave Canada. The depth of assessment is proportionate to the system's risk tier: High-tier systems require a full impact assessment with documented test evidence; Medium-tier systems require a standard assessment; Low-tier systems require a brief assessment. For third-party AI suppliers whose systems may affect customers, we additionally run a structured Third-Party AI Review Procedure that includes a privacy and data governance assessment, a model governance and transparency assessment, and a human oversight verification step. The Privacy Officer confirms whether the supplier will process personal information and verifies that no customer or telemetry data will be used for training or model improvement. These documents are currently at version 1.0 Draft and are not yet in effect.

Source: AI Governance and Accountability Policy, 8. Impact Assessment and Review Before Deployment; Third-Party AI Review Procedure, 6.3 Step 3: Privacy and Data Governance Assessment; Third-Party AI Review Procedure, 6.4 Step 4: Model Governance and Transparency Assessment; Third-Party AI Review Procedure, 6.5 Step 5: Human Oversight and Action Boundary Verification

3. Is personal information used to train AI models?

Partly answered

No. Our policy position is that no customer or telemetry data is used to train, fine-tune or otherwise improve any AI model, whether operated by us or by a third-party provider. This commitment is verified as part of every third-party AI supplier review: the Privacy Officer confirms the supplier's commitment that no data will be used for training, fine-tuning or model improvement, and identifies all subprocessors to confirm each meets the same restriction. This prohibition is also a required contract term for every AI supplier engagement. A privacy impact assessment dated 18 June 2026 for our meter anomaly detection model was reviewed by the Privacy Officer and sets a 24-month retention period for flagged readings with access restricted to billing staff. These policy documents are currently at version 1.0 Draft and are not yet in effect.

Source: AI Governance and Accountability Policy, 8.2 Impact Assessment — confirmation that no customer data will be used for vendor model training; Third-Party AI Review Procedure, 6.3 Step 3: Privacy and Data Governance Assessment — 3.2; Third-Party AI Review Procedure, 6.6 Step 6: Consolidate Findings and Determine Contract Requirements — 6.2; Meter anomaly model privacy review, Privacy impact assessment dated 18 June 2026

To complete this answer: The meter anomaly model privacy review is noted as partially supporting the relevant control. Supplier contracts should be checked to confirm the training-use prohibition is reflected in executed agreements for all current AI suppliers.

4. How often do you review access to AI systems?

Partly answered

We conduct quarterly privileged access reviews. A completed review for the grid telemetry pipeline covering 1 April to 30 June 2026 reviewed 14 privileged accounts: three accounts belonging to departed contractors were disabled on 14 May 2026, two service accounts were found with interactive login enabled (one remediated, one remaining open pending vendor confirmation), and the remainder were confirmed appropriate. The review was performed by the Head of Engineering and recorded in the access register. Our Password and Authentication Standard, which is documented but not yet in effect, requires system owners to conduct access reviews at the frequencies set in that standard and to record results in the access register. The standard also requires that MFA coverage across in-scope accounts is reported to the Chief Operating Officer as part of the annual compliance summary.

Source: Q2 2026 privileged access review, grid telemetry pipeline (claimed A.5.18), Quarterly review of privileged accounts on the grid telemetry pipeline, period 2026-04-01 to 2026-06-30; Password and Authentication Standard, 3. Roles and Responsibilities — System owners; Password and Authentication Standard, 5.6 MFA coverage metric

To complete this answer: No evidence of a quarterly access review for the forecasting service accounts was located for the Q2 2026 period, and one service account remediation on the telemetry pipeline remains open pending vendor confirmation. A completed review record for the forecasting service would be needed to demonstrate full coverage.