

AI Act Readiness Report

Combined executive and auditor report pack for deployer-facing AI Act readiness review.

Company: Example Energy

System: Customer Assistant

Generated at: 2026-03-27T17:09:44+00:00

- Overall status: ready_with_actions
- Risk level: medium
- Bundle fingerprint: 97f83eee6f9f7fdc5c59698fc449b69101d9304ea10961442fc20ecdcb962715

This PDF is derived from the same evidence bundle as the JSON and Markdown outputs. It summarizes readiness findings and reviewer-facing appendix material in one document.

Executive Report

AI Act Executive Report

- Company: Example Energy
- System: Customer Assistant
- Role: provider
- Sector: general
- Overall status: ready_with_actions
- Risk level: medium
- Bundle fingerprint: 97f83eee6f9f7fdc5c59698fc449b69101d9304ea10961442fc20ecdcb962715
- Rulebook version: deployer_readiness_v1
- Rulebook snapshot date: 2026-03-27
- Annex III classification: no_use_cases_declared
- Integrity status: unsigned

Decision

- This bundle is a readiness and evidence pack, not an automated declaration of legal compliance.
- Vauban reports only what it observed, derived, or could not verify in the supplied scope.
- Control outcomes: pass=8, fail=0, unknown=1, not_applicable=12
- Blocking controls: 0; unresolved controls: 1.

Likely Obligations

- Article 4 AI literacy evidence
- Article 5 prohibited-practice screening
- Article 50 transparency for non-obvious AI interaction

Priority Actions

- Attach existing attack/eval reports or run Vauban technical assessments before presenting the readiness pack externally.

Highest-Risk Findings

- control.vauban.readiness.technical_evidence [medium]: No technical testing evidence was attached, so robustness claims would be easy to challenge.

Next action: Attach existing attack/eval reports or run Vauban technical assessments before presenting the readiness pack externally.

Key Gaps

- vauban.readiness.technical_evidence

Risk Register

- Open items: total=1, high=0, medium=1, low=0

Remediation Plan

AI Act Remediation Plan

- Company: Example Energy
- System: Customer Assistant
- Overall status: ready_with_actions

vauban.readiness.technical_evidence

- Status: unknown
- Blocking: False
- Legal review required: False
- Recheckable by Vauban: True
- Required artifacts: At least one technical report artifact
- Missing markers: none
- Action: Attach existing attack/eval reports or run Vauban technical assessments before presenting the readiness pack externally.
- Recheck: Re-run after attaching at least one technical report file.

FRIA Preparation Pack

FRIA Preparation Pack

- Status: not_required
- Required: False
- Control: ai_act.triage.fria_requirement

Trigger Labels

- No FRIA trigger labels were identified.

Evidence Checklist

- Deployment-specific intended purpose and workflow description
- Affected persons and groups, including vulnerable populations
- Provider instructions for use, limitations, and model/service version
- Concrete fundamental-rights, health, and safety risk register
- Human oversight, complaint handling, and redress procedures
- Mitigation measures, monitoring plan, and incident escalation owner

Questionnaire

- Describe the specific context of use, affected persons or groups, and any decisions or recommendations supported by the AI system.
Why it matters: FRIA depends on the concrete deployment context.
- Attach provider instructions for use, intended purpose, limitations, and relevant technical documentation.
Why it matters: Article 27 expects deployers to use provider inputs.
- List the specific risks of harm to fundamental rights, health, or safety for affected persons.
Why it matters: This is the core substantive FRIA input.
- Document human oversight, complaint handling, redress, and escalation arrangements for this deployment.
Why it matters: Mitigations and governance are part of FRIA prep.
- Confirm the internal owner for any authority notification or registration step if FRIA is confirmed as required.
Why it matters: The deployer must operationalize follow-through.

Auditor Appendix

AI Act Auditor Appendix

Scope and Inputs

- Company: Example Energy
- System: Customer Assistant
- Role: provider
- Sector: general
- EU market: yes
- Intended purpose: Answers customer questions with a third-party GPAI service.
- Overall status: ready_with_actions
- Risk level: medium
- Bundle fingerprint: 97f83eee6f9f7fdc5c59698fc449b69101d9304ea10961442fc20ecdcb962715
- Rulebook: deployer_readiness_v1 / snapshot 2026-03-27

Epistemic Policy

- Goal: readiness evidence and gap reporting; not an automated declaration of legal compliance.
- Claim kinds: observed_by_vauban, derived_by_rule, asserted_by_client, requires_external_review
- Completion rule: every applicable control terminates as pass, fail, or unknown.

Control Outcomes

- scope.minimum_facts: pass; applies=True; blocking=True; claim_kind=observed_by_vauban; title=Minimum factual scope for the report
Rationale: Company name, system name, and intended purpose are all present in the config.
Evidence IDs: fact.role, fact.eu_market, fact.intended_purpose
Sources: none
- ai_act.article4.ai_literacy_record: pass; applies=True; blocking=True; claim_kind=observed_by_vauban; title=Article 4 AI literacy evidence
Rationale: Vauban verified that an AI literacy record exists and covers the minimum Article 4 themes pinned in the rulebook.
Evidence IDs: file.ai_literacy_record
Sources: AI Literacy - Questions & Answers (<https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers>); AI talent, skills and literacy (<https://digital-strategy.ec.europa.eu/en/policies/ai-talent-skills-and-literacy>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- ai_act.article5.prohibited_practices_screen: pass; applies=True; blocking=True; claim_kind=derived_by_rule; title=Article 5 prohibited-practices screen
Rationale: No obvious Article 5 prohibited-practice trigger was declared in the current system description.

Evidence IDs: fact.eu_market, fact.uses_emotion_recognition, fact.uses_biometric_categorization, fact.biometric_sensitive_trait_inference, fact.manipulative_or_deceptive_techniques, fact.behavior_distortion_significant_harm, fact.exploits_vulnerabilities, fact.social_scoring, fact.predictive_policing_solely_on_profiling, fact.untargeted_face_scraping, fact.rbi_law_enforcement

Sources: Regulation (EU) 2024/1689

(<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>); Navigating the AI Act

(<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>)

- ai_act.article50.human_interaction_notice: pass; applies=True; blocking=True; claim_kind=observed_by_vauban; title=Article 50 transparency for human interaction

Rationale: A transparency notice exists and contains minimum markers for non-obvious AI interaction disclosure.

Evidence IDs: file.transparency_notice, fact.human_interaction, fact.interaction_obvious

Sources: Guidelines and Code of Practice on Transparent AI Systems

(<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-and-code-practice-transparent-ai-systems>); Navigating the AI

Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689

(<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- ai_act.article50.emotion_biometric_notice: not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 50 transparency for emotion or biometric exposure

Rationale: No emotion-recognition or biometric-categorization exposure was declared.

Sources: Guidelines and Code of Practice on Transparent AI Systems

(<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-and-code-practice-transparent-ai-systems>); Navigating the AI

Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689

(<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- ai_act.article50.deepfake_disclosure: not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 50 deployer disclosure for deepfake or synthetic media

Rationale: No deepfake or synthetic media deployment was declared.

Sources: Guidelines and Code of Practice on Transparent AI Systems

(<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-and-code-practice-transparent-ai-systems>); Commission

launches consultation to develop guidelines and Code of Practice on transparent AI systems

(<https://digital-strategy.ec.europa.eu/en/news/commission-launches-consultation-develop-guidelines-and-code-practice-transparent>

Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- ai_act.article50.public_interest_text_disclosure: not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 50 deployer disclosure for public-interest text

Rationale: No public-interest text publication was declared.

Sources: Guidelines and Code of Practice on Transparent AI Systems

(<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-and-code-practice-transparent-ai-systems>); Commission

launches consultation to develop guidelines and Code of Practice on transparent AI systems

(<https://digital-strategy.ec.europa.eu/en/news/commission-launches-consultation-develop-guidelines-and-code-practice-transparent>

Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- ai_act.triage.high_risk_annex_iii: pass; applies=True; blocking=True; claim_kind=derived_by_rule; title=High-risk and FRIA triage

Rationale: No Annex I product-route or Annex III high-risk trigger was declared in the current system description.

Evidence IDs: fact.role, fact.eu_market, fact.intended_purpose

Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU)

2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- ai_act.triage.article6_3_annex_iii_carve_out: not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 6(3) Annex III carve-out triage

Rationale: No Annex III use case was declared, so Article 6(3) is not in scope.

Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU)

2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- **ai_act.triage.fria_requirement:** pass; applies=True; blocking=True; claim_kind=derived_by_rule; title=FRIA triage
Rationale: No declared public-sector/public-service or Annex III point 5(b)/(c) trigger suggests a mandatory FRIA.
Evidence IDs: fact.role, fact.eu_market, fact.intended_purpose
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.instructions_monitoring:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 26 instructions-of-use and monitoring
Rationale: Article 26 deployer monitoring duties are out of scope unless the report declares a high-risk deployer scenario.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.human_oversight:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 26 human oversight
Rationale: Article 26 human-oversight duties are out of scope unless the report declares a high-risk deployer scenario.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.input_data_governance:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 26 input-data governance
Rationale: Article 26 input-data duties only apply when the declared high-risk deployer provides input data to the system.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.log_retention:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 26 log retention
Rationale: Article 26 log-retention duties are out of scope unless the report declares a high-risk deployer scenario.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.workplace_notice:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 26 workplace information duty
Rationale: Workplace notice duties only apply when the declared high-risk deployer uses the system in the workplace.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.affected_person_notice_and_explanation:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Article 26 affected-person notice and Article 86 explanation readiness
Rationale: Affected-person notice duties only apply when the declared high-risk deployer makes or assists decisions about natural persons.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)
- **ai_act.article26.public_authority_registration:** not_applicable; applies=False; blocking=True; claim_kind=derived_by_rule; title=Public-sector high-risk registration evidence
Rationale: Public-authority registration only applies when the declared high-risk deployer is in a public-sector context outside the critical-infrastructure exception.
Sources: Navigating the AI Act (<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>); Regulation (EU) 2024/1689 (<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>)

- `vauban.readiness.provider_documentation`: `not_applicable`; `applies=False`; `blocking=False`; `claim_kind=derived_by_rule`; `title=Provider documentation available`
Rationale: Provider documentation safeguard only applies to deployers building on third-party GPAI systems.
Sources: none
- `vauban.readiness.technical_evidence`: `unknown`; `applies=True`; `blocking=False`; `claim_kind=observed_by_vauban`; `title=Technical test evidence retained`
Rationale: No technical testing evidence was attached, so robustness claims would be easy to challenge.
Sources: none
- `vauban.readiness.human_oversight`: `pass`; `applies=True`; `blocking=False`; `claim_kind=observed_by_vauban`; `title=Human oversight procedure available`
Rationale: A human oversight procedure exists and a risk owner is named.
Evidence IDs: `file.human_oversight_procedure`
Sources: none
- `vauban.readiness.incident_response`: `pass`; `applies=True`; `blocking=False`; `claim_kind=observed_by_vauban`; `title=Incident response procedure available`
Rationale: An incident response procedure exists and a compliance contact is named.
Evidence IDs: `file.incident_response_procedure`
Sources: none

Evidence Inventory

- `fact.role`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.eu_market`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.intended_purpose`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.annex_iii_use_cases`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.uses_general_purpose_ai`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.human_interaction`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.interaction_obvious`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.emotion_or_biometric_exposure`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.uses_emotion_recognition`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.uses_biometric_categorization`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.biometric_sensitive_trait_inference`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.manipulative_or_deceptive_techniques`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.behavior_distortion_significant_harm`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.exploits_vulnerabilities`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.social_scoring`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.predictive_policing_solely_on_profiling`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.untargeted_face_scraping`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.rbi_law_enforcement`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.rbi_law_enforcement_exception`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.public_interest_text`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.public_interest_text_exception`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.synthetic_media`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`
- `fact.deepfake_creative_context`: `exists=None`; `placeholder=None`; `path=None`; `sha256=None`

- fact.annex_i_product_route: exists=None; placeholder=None; path=None; sha256=None
- fact.article6_3_carve_out_claims: exists=None; placeholder=None; path=None; sha256=None
- fact.workplace_deployment: exists=None; placeholder=None; path=None; sha256=None
- fact.provides_input_data_for_high_risk_system: exists=None; placeholder=None; path=None; sha256=None
- fact.decisions_about_natural_persons: exists=None; placeholder=None; path=None; sha256=None
- fact.legal_or_significant_effects: exists=None; placeholder=None; path=None; sha256=None
- file.ai_literacy_record: exists=True; placeholder=False; path=tmp/pdfs/evidence/ai_literacy.md; sha256=d7c9baa6b303536da158a2066e96d01f5dabafbee7abae539b1cb46baffda7f
Structured fields: last_updated, owner, refresh_cadence, risk_topics, role, scope, system_context, target_roles
- file.transparency_notice: exists=True; placeholder=False; path=tmp/pdfs/evidence/transparency.md; sha256=73f0b0b7efeb7ee7a4d453f02eb0f6ba422e8c78efaa17f1b306562665c0b4
- file.human_oversight_procedure: exists=True; placeholder=False; path=tmp/pdfs/evidence/oversight.md; sha256=64dd22133ac667864fad3700812916671df6ad87f6f17b2c468d07d531a5acfd
Structured fields: escalation_trigger, override_capability, review_step
- file.incident_response_procedure: exists=True; placeholder=False; path=tmp/pdfs/evidence/incident.md; sha256=119df85f01009854bd72b9df5ce17d5dfa782bc62caf2ba217b9334b29fa5444
Structured fields: contact_owner, escalation_or_reporting, incident_scope
- file.provider_documentation: exists=True; placeholder=False; path=tmp/pdfs/evidence/provider.md; sha256=aa9b9c741527789c46a2daded79fc727aa1151c379366a80dfbf25588deb4f82
Structured fields: limitations, model, provider, version
- file.operation_monitoring_procedure: exists=False; placeholder=False; path=None; sha256=None
- file.input_data_governance_procedure: exists=False; placeholder=False; path=None; sha256=None
- file.log_retention_procedure: exists=False; placeholder=False; path=None; sha256=None
- file.employee_or_worker_representative_notice: exists=False; placeholder=False; path=None; sha256=None
- file.affected_person_notice: exists=False; placeholder=False; path=None; sha256=None
- file.explanation_request_procedure: exists=False; placeholder=False; path=None; sha256=None
- file.eu_database_registration_record: exists=False; placeholder=False; path=None; sha256=None

Technical Evidence

- Attached artifacts: 0; existing=0; interpreted_json=0; findings=0
- No technical findings were interpreted from attached artifacts.

Open Risk Register Items

- control.vauban.readiness.technical_evidence [medium]: No technical testing evidence was attached, so robustness claims would be easy to challenge.

Next action: Attach existing attack/eval reports or run Vauban technical assessments before presenting the readiness pack externally.

Sources: none

Integrity and Reproducibility

- Evidence manifest SHA-256: 19b2f91cc684e9bac65721f9cade35070e560981429f5f705e1a5c1cfb69c9ea
- Signature status: unsigned
- Signature algorithm: none
- Signature env var: None

Official Sources

- Navigating the AI Act (European Commission):
<https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>
- AI Literacy - Questions & Answers (European Commission):
<https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers>
- AI talent, skills and literacy (European Commission):
<https://digital-strategy.ec.europa.eu/en/policies/ai-talent-skills-and-literacy>
- Guidelines and Code of Practice on Transparent AI Systems (European Commission):
<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-and-code-practice-transparent-ai-systems>
- Commission launches consultation to develop guidelines and Code of Practice on transparent AI systems (European Commission):
<https://digital-strategy.ec.europa.eu/en/news/commission-launches-consultation-develop-guidelines-and-code-practice-t>
- Regulation (EU) 2024/1689 (EUR-Lex):
<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>