LEARNING · AIGOVOPS FOUNDATION

10 Questions About Your AI Policy

How your organization can create an AI policy, one decision at a time, from the AiGovOps Foundation's open-source work.

Get to Yes · Stay at Yes · Recover to Yes

How to use it. Circle one answer for each question, alone or with your team. The answer key follows the questions, and the last page turns the ten answers into your policy, with a line for who owns each one. The interactive version is at `aigovops-foundation.com/learning/ai-policy-10-questions`.

1

GET TO YES · START WITH WHAT'S REAL

You're writing your organization's first AI policy. Where do you start?

- (A) A statement of the values you want your AI to reflect
- (B) A list of the AI already in use, each with one named owner
- (C) A pause on all AI until the policy is approved

2

GET TO YES · WRITE IT SO IT CAN BE CHECKED

Which sentence belongs in your policy?

- (A) AI-assisted changes to plant software are reviewed by a person and never go straight into the operations network.
- (B) We use AI responsibly and ethically.
- (C) Teams should think about risk before using AI.

3

GET TO YES · ONE RULE, MANY FRAMEWORKS

You answer to the EU AI Act, the NIST AI RMF and ISO/IEC 42001. How should your policy handle all three?

- (A) Write a separate policy for each framework
- (B) Follow whichever framework is shortest
- (C) Write each rule once and map it to every framework it satisfies

4

GET TO YES · A PROMPT IS NOT A RULE

Your support chatbot's instructions say, “Never offer more than a 10% discount.” Is that enough?

- (A) Yes. The model will follow its instructions.
- (B) No. Enforce the limit in a check outside the model.
- (C) Yes, as long as the instruction is repeated several times.

© 2026 AiGovOps Foundation

5

STAY AT YES · THREE ANSWERS, AND NO MAYBE

An AI agent asks to take an action that matters, such as sending money or changing a record. How should your gate answer?

- (A) Yes, no or maybe, and a maybe stays open until someone gets to it
- (B) Yes by default, with a review later
- (C) Yes with evidence, no on the record, or held for a person, and a hold turns into no if nobody approves in time

6

STAY AT YES · THE HOLD

When should a person review an AI decision before it takes effect?

- (A) Before consequential ones, such as moving money or denying a claim, with enough time to actually look
- (B) Only after a customer complains
- (C) Never, because review slows the system down

7

STAY AT YES · WHO APPROVES

Who should approve a new rule in your AI policy?

- (A) The person who asked for it
- (B) Two authorized people other than the requester, with a third person signing the approved version
- (C) The AI that drafted it, since it knows the rule best

8

STAY AT YES · RECEIPTS, NOT REASSURANCE

How do you show that your policy actually ran?

- (A) A slide that says the team is compliant
- (B) The AI vendor's certificate
- (C) A signed receipt for each decision that someone outside your team can verify

9

GET TO YES · VENDOR CLAIMS

A vendor says its AI is “fully compliant.” What does your policy ask for?

- (A) Evidence you can check yourself, before you say yes
- (B) Nothing more. They're the experts.
- (C) A ban on every vendor that uses AI

10

RECOVER TO YES · THE WAY BACK

Your AI system starts causing harm at 9:30 on a Tuesday morning. What should your policy already say?

- (A) Call the vendor and wait for a fix
- (B) Who can turn it off, how, and how you get back to a safe state, practiced in advance
- (C) Let it run until the next review meeting

© 2026 AiGovOps Foundation

ANSWER KEY

The answers, and why

1

B. A list of the AI already in use, each with one named owner

A policy governs real systems. The Foundation's ten-minute Gate Check opens with a scope statement that has an owner's name on it. Values matter, but a rule nobody owns is a rule nobody runs.

From the Foundation: The ten-minute Gate Check · `aigovops-foundation.com/blog/the-ten-minute-gate-check.html`

2

A. AI-assisted changes to plant software are reviewed by a person and never go straight into the operations network.

Write rules someone can check, not a hoped-for outcome. This example comes from the Foundation's Policy Foundry, where every requirement also names its scope, its owner, its risk tier, how exceptions work and when it gets reviewed.

From the Foundation: Policy Foundry · `github.com/aigovops-foundation/aigovops-policy-foundry`

3

C. Write each rule once and map it to every framework it satisfies

One rule, many mappings. The Foundation's Frameworks explorer shows which of forty frameworks apply to you, and Umbrella declares each control once and tracks its coverage across frameworks.

From the Foundation: Frameworks explorer · `aigovops-foundation.com/frameworks.html`

4

B. No. Enforce the limit in a check outside the model.

A polite instruction to a language model is not a gate. In 2023 a dealer's chatbot agreed to sell a new Tahoe for $1 and called it a legally binding offer. Level 200 of the Foundation's practice starts here.

From the Foundation: The one-dollar Tahoe · `aigovops-foundation.com/blog/the-one-dollar-tahoe.html`

5

C. Yes with evidence, no on the record, or held for a person, and a hold turns into no if nobody approves in time

This is the Foundation's gate law: 1 is a yes earned with evidence, 0 is a no written to the record, and ? waits for a human and decays to 0 at its deadline. An open-ended maybe is a no with homework.

From the Foundation: The gate law · `aigovops-foundation.com/resources.html#library`

6

A. Before consequential ones, such as moving money or denying a claim, with enough time to actually look

Level 300 of the Foundation's practice is the hold. ProPublica reported insurance claims denied in seconds each; an Arup employee moved HK$200 million after a deepfake video call. A hold only counts if the review is real.

From the Foundation: Seconds each · `aigovops-foundation.com/blog/seconds-each.html`

© 2026 AiGovOps Foundation

7

B. Two authorized people other than the requester, with a third person signing the approved version

In the Policy Foundry the requester never counts as an approver, two different people approve, a separate signing officer signs, and AI agents can draft and test but never approve or sign.

From the Foundation: Policy Foundry: approvals · `github.com/aigovops-foundation/aigovops-policy-foundry/blob/main/docs/wiki/Policy-and-governance.md`

8

C. A signed receipt for each decision that someone outside your team can verify

Evidence that ships. Beacon writes signed receipts of AI decisions that anyone can check in a browser. The Post Office Horizon scandal shows why a system's own output is not evidence on its own.

From the Foundation: What is Beacon? · `aigovops-foundation.com/beacon.html`

9

A. Evidence you can check yourself, before you say yes

A claim is not evidence. The Gate Check includes a vendor-claim check, and the Foundation's open-source Vendor RFI gives you the questions, a scored self-audit and a receipt you can keep.

From the Foundation: Vendor RFI · `github.com/aigovops-foundation/aigovops-vendor-rfi`

10

B. Who can turn it off, how, and how you get back to a safe state, practiced in advance

Recover to yes. One bad deploy cost Knight Capital $440 million in 45 minutes. The Gate Check asks for the name of the person who last turned the system off, because a kill switch nobody has used is a guess.

From the Foundation: Forty-five minutes at Knight Capital · `aigovops-foundation.com/blog/forty-five-minutes-at-knight.html`

© 2026 AiGovOps Foundation

YOUR AI POLICY

In ten lines

Tick each line when it is true for your organization, and write down who owns it.

- [ ] **1.** We list every AI system we use, and each one has a named owner. Owner
- [ ] **2.** Every rule can be checked, and names its scope, owner, risk tier, exceptions and review date. Owner
- [ ] **3.** We write each rule once and map it to the frameworks we answer to. Owner
- [ ] **4.** Limits are enforced outside the model. A prompt is not a control. Owner
- [ ] **5.** Every consequential action gets a 1, a 0 or a ?, and a ? that nobody answers becomes a 0. Owner
- [ ] **6.** Consequential decisions wait for a real human review. Owner
- [ ] **7.** Two people approve each rule, a third signs it, and AI never approves its own work. Owner
- [ ] **8.** Each decision leaves a signed receipt that an outsider can verify. Owner
- [ ] **9.** Vendors show evidence we can check, not just claims. Owner
- [ ] **10.** We know who can turn each system off, and we rehearse the way back. Owner

One next step: run the ten-minute Gate Check on one AI system you use. `community.aigovops-foundation.com/gate-check.html`

Talk it through: Bob Rapp, `bob.rapp@aigovops.community` · Ken Johnston, `ken.johnston@aigovops.community`

Keep going: the Tuesday Letter at `aigovops.substack.com` and the Library at `aigovops-foundation.com/library.html`

`Created using AI tools`

A starting point drawn from the AiGovOps Foundation's open-source practice, the Policy Foundry and the Library. It isn't legal advice or a certification. © 2026 AiGovOps Foundation. All rights reserved.