

Here's a full editorial blueprint for turning the site into a 9.5-level content product.

## Site map

### Core pages

- Homepage.
- Start Here.
- About / Thesis.
- Contact.
- Newsletter signup.

### Pillar pages

- AI Accountability Architecture.
- Decision Traceability.
- Autonomy and Escalation.
- Governance Operating Model.
- Regulated AI Implementation.

### Supporting article clusters

- Traceability.
- Auditability.
- Oversight and escalation.
- Operational governance.
- Failure handling.
- Industry-specific implementations.

### Asset library

- Templates.
- Checklists.
- Rubrics.
- Decision trees.
- Example logs.
- Framework diagrams.

### Homepage outline

## Hero

One sentence that states the site's promise:

**Practical frameworks for making AI systems auditable, governable, and defensible in production.**

## Immediate navigation

Add three visible choices:

- Start here.
- Read the framework.
- Download a practical asset.

## Featured modules

- Core pillars.
- Latest essays.
- Most useful templates.
- One “recommended path” for new readers.

## CTA block

- Subscribe for new frameworks.
- Download the starter checklist.
- Contact for advisory or collaboration.

## Start Here page

This page should do the onboarding work the homepage cannot.

## Sections

- What this site is about.
- Who it is for.
- The core vocabulary.
- The main framework in plain English.
- Recommended reading order.

## Suggested reading order

1. Thesis page.
2. Traceability pillar.

3. Autonomy and escalation pillar.
4. Governance operating model pillar.
5. A case-study post or implementation guide.

## **Pillar page template**

Every pillar page should have the same structure.

### **Sections**

- Plain-English definition.
- Why it matters.
- Core model or diagram.
- Key decisions teams must make.
- Common failure modes.
- Related posts.
- Downloadable artifact.
- FAQ.

### **Example**

For “Decision Traceability,” the pillar page should explain:

- What counts as evidence.
- Where logs should live.
- Which decisions need traceability.
- How teams prove a decision was defensible.

### **Article template**

Every article should be built to serve readers who need to act, not just learn.

### **Required structure**

- Executive summary.
- Problem.
- Principle.
- Pattern.
- Example.
- Failure mode.
- Implementation steps.

- Checklist.
- Related resources.
- Next action.

## Required ending

Every post should end with one of:

- A checklist.
- A template.
- A decision rule.
- A “what to do next” section.

## Content clusters

### Cluster 1: Traceability

Topics:

- What traceability means in AI systems.
- Audit trails vs logs.
- Evidence chains for automated decisions.
- Replayability and reconstructability.

### Cluster 2: Autonomy

Topics:

- Defining autonomy tiers.
- Human-in-the-loop vs human-on-the-loop.
- Escalation thresholds.
- Override governance.

### Cluster 3: Governance operations

Topics:

- RACI for AI governance.
- Governance board vs operational control.
- Policy-to-process translation.
- Ownership models for AI systems.

### Cluster 4: Failure handling

Topics:

- Model incidents.
- Escalation workflows.
- Drift and anomaly response.
- Audit preparation and response.

## **Cluster 5: Regulated implementations**

Topics:

- Lending.
- Fraud.
- Healthcare.
- Claims.
- Underwriting.
- Customer support.

## **Required asset library**

This is what will move the site from “thought leadership” to “actual toolkit.”

## **Must-have assets**

- AI governance policy template.
- Decision traceability checklist.
- Escalation matrix.
- Autonomy tier rubric.
- Audit log schema.
- Incident response playbook.
- Oversight RACI template.
- Model review intake form.

## **Asset rules**

- Each asset must be linked from at least one article.
- Each asset should be downloadable.
- Each asset should include a short explanation of when to use it.

## **SEO blueprint**

## **Primary keyword themes**

- AI governance.
- AI accountability.
- AI auditability.
- AI traceability.
- Responsible AI operations.

## **Supporting intent themes**

- AI governance checklist.
- AI audit trail template.
- AI escalation process.
- AI decision logging.
- Model oversight framework.

## **SEO execution rules**

- One primary keyword per page.
- One intent per page.
- One pillar page per theme.
- Internal links from all supporting posts to the pillar.
- Supporting posts should answer narrow, practical queries.

## **Editorial voice**

### **Tone**

- Confident.
- Direct.
- Practical.
- Non-hype.
- Built for operators and technical leaders.

### **Style rules**

- Use short paragraphs.
- Prefer concrete examples over abstractions.
- Avoid generic AI commentary.
- State the takeaway early.
- Explain what to do, not just what it means.

### **Visual system**

## Core visual types

- Framework diagrams.
- Control matrices.
- Workflow maps.
- Decision trees.
- Evidence chains.
- Tier models.

## Placement

Use at least one strong visual in:

- Every pillar page.
- Every flagship article.
- Every downloadable asset page.

## Publishing cadence

### Phase 1: Foundation

Publish:

- Homepage rewrite.
- Start Here page.
- 5 pillar pages.
- 5 templates or checklists.

### Phase 2: Depth

Publish:

- 2–3 supporting posts per pillar.
- One applied case study per cluster.
- One visual per major article.

### Phase 3: Expansion

Publish:

- Industry-specific implementations.
- Comparison posts.
- Failure analysis posts.

- Operational playbooks.

## Quality bar

A piece is ready only if it passes these tests:

- A practitioner can use it today.
- A manager can forward it to a team.
- A technical reviewer can see the operational logic.
- A skeptical reader can identify the safeguards.
- A new reader can understand the point quickly.

## Recommended next move

If you want the highest leverage step, I'd build these first:

1. Start Here page.
2. One flagship pillar page.
3. One downloadable checklist.
4. One implementation article with a diagram.
5. One template library page.