For the Abu Dhabi Investment Office

AI maturity documentation workflow

A clear, checked, step-by-step approach

I'll handle the build, the checks and the paperwork, and your team stays in control of every approval.

This follows the Department of Government Enablement framework already used at ADIO: seven pillars, 81 capabilities and five ratings. Finishing a document does not, by itself, set a maturity rating. Humans keep approvals, formal appointments, committee activity and operating evidence. Agents may propose responsibility assignments but cannot confer authority.

How the work flows



Your team of 7 AI agents, each with one job

01 **Manager**
Plans and assigns the work.

02 **Knowledge custodian**
Keeps the sources and page references.

03 **Document architect**
Plans the smallest set of documents.

04 **Writers**
Draft each document.

05 **Assessment reviewer**
Checks against the rating.

06 **Evidence verifier**
Checks every claim has a source.

07 **Editor & publisher**
Finalises and versions.

Each is a separate Codex agent. None of them can approve. Neither the manager nor a writer may override a failed review.

What I'll do for you, step by step

01
You confirm first
Before any drafting starts, you confirm the source versions, the scope, the target ratings and the dates, and resolve differences between the uplift plan and the evidence schedule.

02
Gather your rulebook
Collect the government handbooks and your plans in one place.

03
Map what's needed
List every requirement for each of the 81 capabilities.

04
AI writes the documents
AI drafts each policy, procedure and template to match the handbooks.

05
AI double-checks
Two separate AI reviewers check every draft. Both checks confirm the edited version. A rejection reopens the affected work.

06
You approve
Nothing is final until your team signs off.

07
Track progress
A simple dashboard shows what's done, what's waiting, and what needs your team.

Any merge or edit invalidates prior passes, and the affected documents are re-reviewed together. Work is also checked again if a source changes.

We start small: AI Governance first, proven on your own tests

The pilot is proven when each of these deliberate test cases stops a document from being released:

01 A deliberately missing requirement
02 A false approval
03 A conflicting role assignment

We also show recovery after a stop, that a run cannot start twice by mistake, and that work is checked again if a source changes.

To get started, I'll just need

- ✓ DGE Handbooks v1.2
- ✓ Reworked uplift plan
- ✓ Evidence taxonomy
- ✓ Oct–Nov evidence schedule
- ✓ In-scope capabilities and target ratings
- ✓ Access to approved hosting
- ✓ Existing ADIO policies and approved records

Ali Mohammadi
ali.m.dataconsulting@gmail.com · +971 58 584 3732

How it runs on Codex

Codex is the AI coding and automation tool that does the work. Here's how the pieces connect, and where your team stays in control.



The manager runs many documents at once and hands each step to the right agent.

Inputs are ranked by authority. The handbooks' rating descriptors govern; indicative evidence examples are not mandatory document lists. Supplementary guidance cannot silently override the handbooks. Conflicts are escalated, not silently aligned.

SAFETY RAIL · SPEND LIMIT · CONCURRENCY LIMIT · AUDIT LOG · STOP BUTTON · NO DUPLICATE RUNS · DAILY REPORT

LEGEND — Live path · Fail loop · Human step · Safety rail

ALI MOHAMMADI
ali.m.dataconsulting@gmail.com · +971 58 584 3732

For the Abu Dhabi Investment Office

Deliverables, acceptance and operation

What I'll hand over, how we test it, and how it runs day to day

Deliverables

01
Configurations and role instructions
The Codex setup, and the instructions each agent follows.

02
Source register
A filed record of sources, with versions and page references.

03
Requirements matrix
Every requirement, with the fields and coverage shown in the store.

04
Editable documents
The drafts and approved documents, in a form your team can edit.

05
Review logs
A record of every assessment and evidence check, including failures.

06
Status dashboard
What's done, what's waiting, and what needs your team.

07
Deployment and recovery runbook
How to run it, how to stop it, and how to recover after an interruption.

Acceptance

01
Coverage
Every in-scope criterion has a source-linked output or an explicit human dependency. No missing requirements are marked complete.

02
Quality
Both reviewers pass released drafts. There are no unresolved critical findings, contradictions or unsupported claims.

03
Reliability
We demonstrate recovery, duplicate-run prevention, and re-review after a source change.

The three tests. Each must block release:

01 A deliberately missing requirement

02 A false approval

03 A conflicting role assignment

Operation

Work runs in scheduled batches around the clock, and pauses when it is complete or blocked. It stays within spending and concurrency limits. After an escalation, independent tasks continue.

Tasks, decisions and checkpoints are saved so work can resume after a stop. There is one writer per document. Progress and blockers are reported each day.

Humans keep approvals, formal appointments, committee activity and operating evidence. Agents may propose responsibility assignments but cannot confer authority.

How we can work together

Whatever scope you need, from one pillar to all seven, I can deliver it.

01
Start with a pilot
We prove it on AI Governance against your own tests, then you decide.

02
Full build
All 7 pillars done, with everything handed over (documents, settings, dashboard and a runbook).

03
Build and run
I build it and then run it day to day, handling batches, questions and daily updates.

Ali Mohammadi

ali.m.dataconsulting@gmail.com · +971 58 584 3732