

SIX-WEEK PROGRAMME FOR TECHNICAL TEAMS

# Secure AI in your repository: a 6-week programme

Help the team develop with AI, with controls you can measure before and after

Six weeks that cover a diagnosis of your configurations, the full workshop, a management report with five prioritised actions and a second measurement using your own numbers. I won't promise you a percentage improvement. What changes, I leave measured.

---

6

weeks, from diagnosis to review

2

measurements: the baseline and the one in week 6

5

phases, each with a deliverable

90 min

remote review in week 6

## THE PROBLEM YOU ALREADY HAVE

# One day of training changes what the team knows, but not always what it does.

After a workshop everyone goes back to their branch, their deadline and their habits. If nobody looks at what got applied, three weeks later you have a nice memory and very few changed configurations. With AI it's worse: assistants, MCP servers and autonomy modes change every few months, and what you decided yesterday goes stale.

There is a second problem: **feeling faster is not measuring**. In 2025 a randomised trial by METR with experienced developers found they took longer with AI while believing they had gone faster. So this programme measures at the start and at the end, and I don't tell you how it should go.

The programme puts together three things that are usually bought separately: the training day, the follow-up afterwards, and the evidence an auditor asks for. **At the end you have changes in your repository, a report to show management and two measurements to compare.**

## Five phases in six weeks

### 01 Diagnosis and baseline

Week 0. A team survey, a review of your AI configurations and the first indicators.

### 02 Workshop and evidence

Week 1. A 5 to 8 hour workshop depending on the team, with labs adapted to your stack, plus the evidence pack.

### 03 Management report

Week 2. Three to five pages with five prioritised actions and an owner for each.

### 04 Questions and pull request review

Weeks 1 to 5. A channel to ask me things, with a closing date, and written review of two or three real pull requests.

### 05 Review and second measurement

Week 6. Ninety minutes remotely and the before-and-after comparison.

A one-day workshop is forgotten within a week. This programme keeps going for six.

## WHAT CHANGES AT THE END

# Five outcomes you can show whoever decides

Each outcome has a concrete deliverable and a week in which it appears. What I can't promise is how much the indicators will move. Your own data tells you that.

01

### Know where you are starting from

With a ten-minute team survey and a review of the AI configurations already in your repositories, we get in writing what the assistant reads, what it is connected to and what it can run.

**Evidence: a diagnosis sheet and a baseline measurement with 4 to 6 agreed indicators.**

02

### Leave the workshop with controls in the repository

In the last block each participant applies the controls to your code. The configurations stay on a branch of your repository and your team merges it.

**Evidence: a branch with configurations, a technical AI and MCP usage policy, and a Definition of Done.**

---

03

### Have evidence for the control, and more than the memory of a course

Signed attendance, an individual certificate and a syllabus mapped to the standards that apply to you, with a one-page letter for whoever audits.

**Evidence: an evidence pack delivered in the week of the workshop.**

---

04

### Give management a document it can read

What was seen, what was worked on, which risks showed up in your configurations and five actions with an owner. It is not an audit and the report says so on its first page.

**Evidence: a 3 to 5 page report in week 2.**

---

05

### Check what was really applied

At six weeks we repeat the indicators and go through with the team what stuck, what was dropped and what needs adjusting.

**Evidence: a before-and-after comparison with your own numbers.**

# What happens each week and who does it

I do most of the work without the team having to stop. I ask for your time in two places: the survey, ten minutes per person, and the day of the workshop.

---

## Week 0

### Diagnosis and baseline measurement

The team answers a ten-minute survey. I look at the AI configurations already in your repositories and we agree four to six indicators. With that I adapt the labs to your stack.

**Deliverable: Diagnosis sheet and baseline measurement.**

---

## Week 1

### Workshop and evidence pack

I give the full workshop, 5 to 8 hours depending on the team, with the last block on your repository. Each participant signs the attendance sheet and when the day is over I prepare the course evidence.

**Deliverable: A branch with configurations, a technical AI and MCP usage policy, and the evidence pack.**

---

## Week 2

### Management report

I write a three-to-five-page document for whoever decides, with what the diagnosis showed, what we worked on in the workshop and five prioritised actions with a proposed owner for each. It is not an audit.

**Deliverable: A 3 to 5 page report.**

---

## Weeks 1 to 5

### Q&A channel and pull request review

We open a channel in your Slack or Teams and I reply within two working days. The team sends me two or three real pull requests and I comment on them in writing, on your platform with your permissions, or with snippets if you prefer. The channel closes at the end of week 5.

**Deliverable: Written review of two or three pull requests.**

---

## Week 6

### Remote review and second measurement

Ninety minutes remotely with the team to see what was applied, what was not and why. We repeat the indicators and adjust any configuration that needs it.

**Deliverable: Before-and-after comparison.**

Week 1 is the full workshop. Its syllabus, labs and agenda are at [danielalfocea.com/en/formacion/workshop-desarrollo-seguro-con-ia](https://danielalfocea.com/en/formacion/workshop-desarrollo-seguro-con-ia)

## How to tell if it helped

In 2025 a randomised trial by METR with experienced open-source developers found they took 19% longer with AI, while believing they had gone 20% faster. Those were early-2025 tools and it says nothing about your team, but it shows that impressions are no measurement. So I won't promise you a percentage. We pick together the indicators that are useful for your team, and these are candidates.

### **Pull requests with generated code reviewed against the checklist**

What share of assistant-made changes goes through the agreed review checklist.

### **SAST and secret findings per pull request**

What your tools already detect, counted before and after on the same kind of change.

### **Review time for a generated pull request**

How long a person takes to accept or reject, measured in your version control system.

### **Assistant configuration changes that get reviewed**

How many edits to files such as mcp.json or IDE settings receive human review.

### **MCP and Skills inventory with an owner and a reason**

How many of the connected ones have someone who answers for them.

### **Team confidence in generated code**

A short survey, repeated in week 6. It is a perception and I treat it as one.

# Course or programme

| IF YOU NEED                                                                 | WORKSHOP                                                                                               | PROGRAMME                                            |
|-----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| <b>The team learns to use assistants under controls</b>                     | The course does it: between 5 and 8 hours depending on the team, with the last block on your own code. | Adds a diagnosis up front and a review afterwards.   |
| <b>Documented justification for the training for a standard or an audit</b> | The course includes the evidence pack.                                                                 | Adds the management report and the final comparison. |
| <b>Knowing what was applied and what changed after a few weeks</b>          | The course does not cover it. That is up to you.                                                       | Two measurements and a review at six weeks.          |
| <b>A team that will apply the controls alone and has the time to do it</b>  | That is all you need. If it gets stuck, the 30-day guarantee covers an hour.                           | Not needed.                                          |

## WHAT STAYS IN THE COMPANY

### Six weeks later, this is what you're left with

On top of all the workshop material, the programme leaves documents made to be passed around inside the company.

#### Diagnosis sheet

How your AI configurations look at the start, and what the team says in the survey.

#### Management report

Three to five pages with five prioritised actions and who answers for each.

#### Written review of two or three pull requests

Comments on your own real changes, using the workshop's checklist.

#### Baseline and final measurement

A template with the indicators we choose together, filled in week 0 and again in week 6.

#### Evidence pack

Signed attendance, individual certificates, a mapped syllabus and a one-page letter for the auditor.

#### Before-and-after comparison

What was measured, what changed and what did not, as it is.

WHAT IS INCLUDED

# Everything in the course, and the rest of the programme on top

The programme includes the workshop and the four bonuses that come with it. If you need to explain what you're getting, start with this list.

- +1

**Compliance evidence pack**  
Signed attendance, an individual certificate with date, hours and contents, a syllabus mapped to the standards that apply to you, and a one-page letter for the auditor.
- +2

**Ten-minute survey for the team beforehand**  
It helps me adapt the labs to your stack and is the basis of the baseline measurement.
- +3

**Configurations on a branch of your repository**  
I leave them in the final block, together with a technical template for an AI and MCP usage policy. It is technical and not legal, so your legal team adapts it.
- +4

**Labs repository with solutions**  
So you can repeat them with whoever joins, under an internal-use licence for 12 months.

COMPLIANCE

# What evidence the programme leaves and which controls it supports

The pack serves as evidence for these controls, but it does not satisfy them for you. The company runs and maintains the control, and this documents the training. When we talk I tell you which ones apply to your case.

| STANDARD                           | WHAT IT ASKS FOR                                                                                                                          | WHAT I LEAVE AS EVIDENCE                                                       |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| PCI DSS 4.0 · 6.2.2                | Documented training, at least every 12 months, for people who develop bespoke software, on security relevant to their role and languages. | A syllabus mapped to the requirement's points, signed attendance and date.     |
| ISO 27001:2022 · A.6.3             | Security awareness, education and training appropriate to each person's role.                                                             | A record per person and role, with contents and date.                          |
| ISO 27001:2022 · A.8.25 and A.8.28 | A secure development life cycle and secure coding, including training for developers.                                                     | A mapped syllabus, plus the resulting Definition of Done and review checklist. |
| SOC 2 · CC1.4 and CC2.2            | A documented training programme, with completion records per person and depth according to role.                                          | A list with name, date, content and role, and an individual certificate.       |

| STANDARD                 | WHAT IT ASKS FOR                                                                                                                             | WHAT I LEAVE AS EVIDENCE                                                                |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| <b>DORA · art. 13(6)</b> | Digital operational resilience training proportionate to the role, for financial entities subject to the regulation.                         | A training record and syllabus, for your file.                                          |
| <b>AI Act · art. 4</b>   | Measures to support AI literacy among staff, adapted to their role and context. It is an obligation of means and no certificate is required. | An internal record of the training, with content on AI risks and use according to role. |

#### WHO IT IS FOR

## For teams that want to see what changed and not just say they attended

### A strong fit for

- Security architecture, AppSec and Platform Engineering owners who need to show results to management.
- Companies under PCI DSS, ISO 27001, SOC 2 or DORA where training already has to be documented.
- Engineering managers who see generated code in production and want to know where their team stands.
- Teams who use Claude Code or GitHub Copilot and decide permissions case by case.

### Not the right programme if

- You want an inspirational talk for the whole department.
- You want to train two hundred people. A training platform will cost you less.
- The team has never used an assistant on a real repository.
- You expect an audit of your code with a vulnerability report, or something that certifies you comply with a standard.

## WHO TEACHES IT

# Security and development from someone who has worked on both sides



I'm Daniel Alfocea, and I've spent twenty years in development and security. At Banco Santander I was Principal Engineer in the AI department when the area was being set up. At BBVA I was Cybersecurity Tech Lead. At 42Crunch, in the UK, I was responsible for one of its API security products.

I've presented research at RSA Conference and OWASP AppSec EU, and I'm a regular speaker at RootedCON. Spain's National Cybersecurity Institute (INCIBE) has asked me to judge its competitions several times. I founded Navaja Negra and the OWASP Madrid chapter. I'm the author of widely used security tools such as GoLismero, Plecost and Dockerscan, and of LAPIS, research published on arXiv on how to describe an API to a model with fewer tokens. With Alfonso Muñoz I wrote *Secure MCP: A practical guide for developers, software architects and tech leads*.

I run this programme myself from start to finish, from the diagnosis to the week 6 review. I review the pull requests too, and I don't hand them to anyone else.

## GUARANTEES

# What I guarantee, in writing

### **If the configurations do not work, I fix them**

If the configurations I leave in your repository do not work, I run a free remote session to correct them.

### **If you are still stuck at 30 days, one more hour**

If after 30 days the team has got stuck applying something from the programme, I give you one hour remotely, for free.

This is not an audit and I don't promise percentages, neither fewer vulnerabilities nor more productivity. What I leave measured is your before and your after.

## Questions worth settling now

---

### How is it different from the workshop?

The workshop is the in-person day, 5 to 8 hours depending on the team, and it is part of the programme. What you add is what happens before and after: the diagnosis, the measurement, the report for management, the follow-up and the week 6 review.

---

### How much of our time does it take?

Ten minutes per person for the survey, the day of the workshop, the ninety-minute review in week 6, and whatever time you choose to spend sending me pull requests and questions. Nothing else.

---

### Do you review our code?

I review two or three pull requests the team sends me, and the last block of the workshop is done on your repository. All of it on your platform with your permissions, or with snippets if you prefer. It is not an audit and there is no vulnerability report.

---

### What if the indicators do not improve?

I tell you anyway. I measure to find out what happened, not to prove it worked. If something did not move, in the week 6 review we look at why and what to adjust.

---

### How much does it cost?

It depends on the duration and the size of the team. After a first conversation I send you a fixed proposal.

---

### Does it include a certificate?

Each participant receives an individual certificate with date, hours and contents. There is no exam and it is not an official accreditation.

---

### Which standards does the evidence pack cover?

It serves as training evidence for PCI DSS 4.0 (6.2.2), ISO 27001, SOC 2, DORA and article 4 of the AI Act. It does not certify that you comply with any of them, because the company runs the control. Before we start I tell you which make sense in your case.

---

### Can it be subsidised through FUNDAE?

Ask me when you request a date. I will send whatever your accountancy firm or training department needs to assess it.

---

### Does it guarantee we will have no AI incident?

No. What I guarantee is the phases described, the delivery of the documents and the two guarantees above. Lower risk depends on the company applying and maintaining the controls afterwards.

## FORMAT

# One team, six weeks and a fixed proposal

The price depends on the duration and the size of the team, and we fix it after a call, based on your stack, the number of repositories and the indicators you choose. It includes the diagnosis, the workshop, the evidence pack, the report, the follow-up and the second measurement.

- Participants from one team or collaborating functions.
- Diagnosis, workshop, evidence pack, report, follow-up and final review.
- In English or Spanish.
- The company provides room, screen and wifi on the workshop day. Everything else is remote.
- Travel outside Madrid is agreed before the date is fixed.

Write to me at

[me@danielalfocea.com](mailto:me@danielalfocea.com)

On the call we see if we fit, what dates work and which indicators make sense to measure in your team.

Can be subsidised through FUNDAE (Spanish training credit).

## EMAIL

[me@danielalfocea.com](mailto:me@danielalfocea.com)

## PROGRAMME ONLINE

[danielalfocea.com/en/formacion/programa-desarrollo-seguro-con-ia](https://danielalfocea.com/en/formacion/programa-desarrollo-seguro-con-ia)