WHAT'S INSIDE

Validating AI in GxP

A Practitioner's Guide · contents and what to expect

This is a working handbook for validating AI and machine learning in GxP environments. It assumes you already know CSV and GxP, and it begins where that training stops: what changes once a system can learn, how to classify and govern it, how to build evidence that holds, and how to keep it valid long after go-live. Read it front to back, or open it at the chapter you need on a Monday morning. Here is what's inside.

Getting your bearings (Chapters 1 to 3)

The concepts that matter for validation, the regulatory ground you're standing on (Annex 11, Draft Annex 22, 21 CFR Part 11 and FDA CSA), and a clear account of where traditional CSV stops working once a model can learn. Start here if AI is new ground for you.

Classify and govern (Chapters 4 to 6)

The four-tier risk model for deciding how hard to test, the governance operating model that gives AI an owner and a board, and the AI Validation Master Plan that ties the programme together. This is the spine of the method.

Build the evidence (Chapters 7 to 10)

Data governance for AI, performance acceptance criteria you can defend, test data and the three independences, and IQ, OQ and PQ rewritten for systems that learn. The chapters that turn intent into an evidence package.

Keep it valid over time (Chapters 11 to 14)

Ongoing monitoring and drift response, explainability and confidence, change control and retraining, and risk management beyond the tier model. Validation for AI carries on past release, and this is where you make it last.

Prove it and scale it (Chapters 15 to 19)

Inspection readiness, generative AI in non-critical use, large language models in regulated work, the tools you can use today, and a first-100-days plan for standing the programme up. How to face an inspector, and how to begin on Monday.

The toolkit (Appendices A to O)

The part you'll reach for most: 28 template skeletons (T01 to T28), 6 SOP skeletons, 11 quick reference guides, a glossary, an inspection question bank, common findings and how to avoid them, a vendor assessment framework, and a worked evidence pack with an Annex 22 crosswalk. Built to be adapted to your own quality system.

Full contents

1 AI in GxP: the concepts that matter for validation

2 The regulatory spine: Annex 11, Draft Annex 22, Part 11 and CSA

3 Why AI validation breaks traditional CSV

4 Classification and risk tiering

5 The AI governance operating model

6 The AI Validation Master Plan

7 Data governance for AI

8 Performance acceptance criteria

9 Test data and the three independences

10 IQ, OQ and PQ redefined for AI/ML

11 Ongoing monitoring and drift response

12 Explainability and confidence

13 Change control and retraining

14 AI risk management beyond the tier model

15 Inspection readiness

16 Generative AI in non-critical GxP use

17 Large Language Models in regulated environments

18 AI in practice: tools you can use today

19 Building the AI programme: the first 100 days

Appendices

A Template skeletons (28) · B SOP skeletons (6) · C AI in Software as a Medical Device · D CDMO and outsourced GMP validation · E Quick Reference Guides (11) · F Glossary · G Template cross-reference index · H References · I Inspection question bank · J Common findings and how to avoid them · K Vendor assessment framework · L Worked evidence pack and Annex 22 crosswalk · M Full reference templates · N Integrating the AI lifecycle into an existing QMS · O Where clinical-trial AI sits, and where this handbook stops