Dr FC Practice Standard

Good AI Practice in Clinical Research

For those who build, provide or use AI

1.0 · 26 November 2026

Dr. Christoph Franz

Dr. Franz Clinical Consulting GmbH, Heddesheim, Germany

Version 1.0, first edition. Open to comment.

Foreword, at a glance, sections 1 to 6, annexes A to C, version history

Foreword

I have worked in clinical development for more than a decade. In that time I have seen good tools fail in teams that had not agreed on how they work, and simple tools succeed in teams that had. Artificial intelligence does not change this. It makes it more visible, and faster.

This document is my answer to a question I have long asked myself, and that I keep meeting in conversations with people in biotech and pharma and in the public discussion: what does good practice look like when AI enters clinical research? It is not a regulation and it does not replace one. It describes the practice I aim for in my own work. It draws on my experience in clinical development and in building AI services. I offer it as a proposal to everyone working in this field: it brings together what regulations and guidelines already say with that experience, and it is meant to be discussed and improved.

I wrote it for everyone who builds, provides or uses AI in clinical research: sponsors, investigators and their institutions, CROs and other service providers, quality leads and clinical operations. If you disagree with a point, I would like to hear it. A standard that is not challenged does not improve. Comments are welcome at info@drfranzclinical.com.

Dr. Christoph Franz
Heddesheim, November 2026

At a glance

This summary is not part of the standard; the full text applies. The document sets out good practice for anyone who builds, provides or uses artificial intelligence (AI) to support clinical research. It is not a regulation and has not been issued or endorsed by any authority. It is its author's proposal, built on existing guidelines and on practical experience, and open to comment. It is modelled on the structure and style of ICH E6(R3) Good Clinical Practice.

Its core ideas are:

1. **People decide.** AI output informs a qualified individual, who remains responsible for any decision taken on it.
2. **Know the use.** Before a tool is used, write down what it is for, what goes in, what comes out and who relies on it. Whether the output becomes part of a trial record changes which requirements apply. What this guidance recommends differs with the task: a tool that helps a team find its way needs less than one that drafts a trial record or supports a decision about participants (examples in Annex A).
3. **Match effort to risk.** The more an error could affect participants, trial results or confidential information, the more testing, checking and documentation are recommended.
4. **Fit for purpose.** This guidance treats an AI service used in a trial as a computerised system. The sponsor or investigator responsible for the trial activity remains responsible for its fitness for purpose, also when it is built by others.
5. **Protect data.** Agree confidentiality and data protection terms first, use only the data needed, and know where data are and who can see them.
6. **Control change.** Know which version produced a result, and test changes before use.
7. **Measure, do not assume.** AI can answer the same question differently. Measure how much, and keep monitoring. Measurement supports the judgement of the person who decides (idea 1); it does not replace it.
8. **Show evidence and limits.** Each statement should point to what it rests on, and the output should say what the tool could not assess.
9. **Keep records and learn.** Keep records that let a result be traced, and correct errors at their root cause.

Eight of these nine ideas bundle the twelve principles of section 2: data protection covers Principles 4 and 5, evidence and limits cover Principles 8, 9 and 11, records and learning cover Principles 10 and 12. Idea 4 rests on 5.5.6 and 5.5.7. Sections 4 and 5 say who does what. Annex A gives examples of use, and Annex B shows where ICH E6(R3) addresses related topics.

If you only use a tool that others built, seven things matter most. The list does not replace the full text.

1. Write down what you use the tool for and whether its output can reach a trial record (5.1).
2. Decide which data may go into it and which may not (5.2).
3. Obtain the provider's evidence and check the settings you rely on (5.10).
4. Test it on your own examples before you rely on it (5.5).
5. Have a qualified person check the output before it is used (5.7).
6. Keep a record of what was used wherever output supports a trial record (5.8).
7. Record errors and correct their causes (5.12).

READING PATHS · 1.4 ⛶ Overview ☐ Use a tool that others built ⊞ Build, operate or assess

1 Introduction

Overview

Use a tool that others built

Build, operate or assess

This guidance sets out how AI services that support clinical research should be defined, built, used, changed and monitored, whoever builds or uses them: sponsors, investigators and their institutions, contract research organisations (CROs) and other service providers. It draws on practical experience of building and using AI services.

1.1 Purpose

AI can draft, summarise, check and search clinical research material at a speed and scale that people working alone cannot, wherever it supports clinical research: in trial design and planning, in documents and submissions, in data review, and in quality and safety work. AI tools can also be wrong in ways that look convincing, vary between runs, and handle confidential information in places nobody checked. The purpose of this guidance is to describe a practice in which the benefit is kept and those risks are controlled, documented and stated honestly to the people who rely on the output.

1.2 Scope

This guidance applies to AI tools, applications and services that support clinical research: for example the design and review of protocols and other trial documents, planning and feasibility, data review, and quality and safety work. It covers the whole life of a service: defining its intended use, handling the data and documents entrusted to it, building and changing the service, measuring how it performs, releasing results, keeping records, and learning from errors. In this guidance, "service" stands for any such tool, application or service, whether it is bought, built in-house or provided to others.

It is written first for uses in which the AI output informs a qualified individual who makes the decision. Where output is adopted into a trial record or carries a decision in a GxP process, this guidance continues to apply, and the responsible party's requirements for computerised systems and essential records apply in addition (see Principle 2, 5.1 and 5.5.6). Whether that is the case is decided by the responsible party, not by this guidance.

It does not cover AI used in direct patient care, in medical devices, or in manufacturing. The principles in section 2 may be useful there as well, but those fields have their own requirements, and this guidance was not written for them.

AI that interacts directly with trial participants, for example in recruitment or in the informed consent process, raises further questions, among them ethics committee review and the requirements for informed consent (ICH E6(R3), Annex 1, section 2.8). This edition does not address them. The principles apply to such uses but are not sufficient for them.

1.3 Relationship to other guidance

This document is modelled on the structure and style of ICH E6(R3) Good Clinical Practice: principles first, then the detail that puts them into practice. It does not replace any regulation or guideline, and it does not claim conformity with any of them. It has not been issued, reviewed or endorsed by ICH, EMA, FDA or any other authority, and it was not developed through a consensus process. "Standard" describes its purpose, not a formal status. It draws in particular on:

- ICH E6(R3), whose principles call for quality built into design, proportionate risk-based approaches, clear roles and reliable results;
- the joint EMA and FDA Guiding Principles of Good AI Practice in Drug Development (January 2026), which this guidance takes into account; the table after section 2 shows where each is addressed and where coverage is partial;
- ICH Q9(R1) on quality risk management;
- ICH M15 on general principles for model-informed drug development, which lists artificial intelligence and machine learning among modelling approaches; it is relevant where AI-based models generate evidence for regulatory decisions, a use beyond the main focus of this guidance;
- the EMA Guideline on computerised systems and electronic data in clinical trials, including its data integrity expectations (ALCOA++), and, on data integrity, the MHRA GXP guidance and the PIC/S guidance for GMP/GDP environments;
- the EU General Data Protection Regulation and the EU Artificial Intelligence Act as amended in 2026.

Full references, with version and date, are listed in section 6.

1.4 How to read this guidance

Section 2 states twelve principles. Each is followed by a short rationale. Section 3 defines terms. Section 4 sets out roles. Section 5 gives lifecycle requirements, grouped by stage.

The wording follows ICH usage. **"Should"** marks a recommendation for anyone who builds, provides or uses an AI service in clinical research. The start of section 5 shows which requirements usually fall to whom.

Not every reader needs every section:

If you want	Read
⛶ An overview	At a glance, and the principle statements in section 2
☐ To use a tool that others built	At a glance with its seven points; 4.5; the right-hand column of the table at the start of section 5; Annex A
⊞ To build, operate or assess a service	The full text; Annex B for the link to ICH E6(R3)

2 Principles of good AI practice

Overview

Twelve principles should guide anyone who builds, provides or uses AI to support clinical research decisions. Principles 1 to 4, 11 and 12 apply to every use. Principles 5 to 10 describe what a service itself should do. Whoever builds or operates the service applies them directly. An organisation that uses a service built by others, for example an office assistant, meets them through its choice of service, its settings and rules and the provider's evidence, as far as the service allows. That is usually a task of the functions that select and configure the tool, such as IT and quality, not of the individual user (see the table at the start of section 5 and 4.5).

Principle 1. Accountable people make the decisions.

AI output in clinical research should inform a qualified individual, who remains responsible for the decision taken on it. Every result that is relied on should be traceable to a named role that checked it.

Rationale.

An AI tool cannot be held accountable, and a reader cannot judge a result whose owner is unknown. Keeping the decision with a qualified individual also reduces the influence of the tool on the decision, which risk-based frameworks take into account.

Principle 2. The intended use and context of use are defined before the service is used.

Each service should have a written intended use that states what it takes as input, what it produces, who uses the output and for what. The decisive question is whether the output will ever become a GxP record or carry a regulated decision.

Rationale.

The same tool can be a low-risk aid or part of a regulated process, depending only on how its output is used. A clear statement of use lets the responsible party and the user agree which requirements apply, and shows when a change of use changes the requirements.

Principle 3. Controls are proportionate to risk.

The depth of testing, checking and documentation should follow the factors critical to the quality of the activity the service supports and the potential impact of an error on participant safety, on the reliability of trial results and on the confidentiality of entrusted information.

Rationale.

ICH E6(R3) and ICH Q9(R1) both call for effort to follow risk. Uniform maximum control wastes effort on trivial tasks and leaves less for the ones that matter.

Principle 4. Confidentiality and data protection are designed in.

Entrusted data and documents should be handled under agreed confidentiality terms before they are received: between organisations, a confidentiality agreement and, where personal data may be present, a data processing agreement; inside one organisation, the access and data protection rules that apply there. Data should be minimised, kept within a named circle of people and systems, and kept or deleted according to a defined schedule that respects any retention the responsible party is required to apply.

Rationale.

Unpublished trial documents and data are among the most sensitive assets in clinical research. Trust lost through a data incident is rarely regained, and data protection law applies wherever personal data are processed, whatever the GxP status of the output.

Principle 5. Only inputs the service can handle safely are accepted.

A service should be set up to work within its safeguards: it should say which kinds of input it is meant for, and refuse or mark inputs for which its safeguards have not been shown to work. How this is done depends on the purpose: a service that takes one document can check it at intake, while one that draws on many sources applies rules for admitting them (see 5.4.5).

Rationale.

A safeguard that has only been tested on one type of input gives no assurance on another. Refusing an unsuitable input is cheap; processing it without protection is not.

Principle 6. Every build is identified, controlled and tested before release.

The configuration that produced a result should be fixed and identifiable. Any change that could alter results should pass regression testing against reference material before it is used in production.

Rationale.

AI services change often, and small changes can shift results. Without a fixed, identified configuration, neither the responsible party nor the user can say which version produced a statement or whether a later result is comparable.

Principle 7. Performance is measured, not assumed.

The performance of a service, including how much its results vary between runs of the same input, should be measured on reference material and monitored over time for drift against defined tolerances. A service should be described as fit for a use only on the basis of measurement.

Rationale.

Many AI systems, in particular those based on language models, do not return identical output on every run. The honest response is to measure the variation, set a tolerance, and look at what moves, rather than promise determinism or ignore the issue. The absence of complaints is not evidence of quality.

Principle 8. Every output carries its evidence.

Whatever an AI service puts in front of a person to rely on (an assertion, a figure, a summary, a recommendation, a draft; in this guidance a "statement") should point to the input it rests on, precisely enough to be checked, and, where a standard, guideline or rule is applied, name it. For a document that is the location and a verbatim quote; for data, the records and the calculation; for a generated text, the sources used. What cannot be supported in this way should be withheld or clearly marked as unsupported, and, where the service selects what it shows, the number withheld should be reported.

Rationale.

An expert should be able to check any such statement in minutes without redoing the work. Reporting how much was withheld shows that selection took place and keeps the output honest about what the tool produced.

Principle 9. Limits are stated, not hidden.

An output should say what it covers and what it does not: which inputs were used and how fully, what the service could not do or could not judge, and where a result is uncertain. A negative result about the subject ("not met", "not found") should be kept apart from the service's own inability to answer ("not assessable"), and the output should say which applies.

Rationale.

A missing answer can mean that something is not there or that the tool could not see it. Mixing the two misleads the reader in one of two directions: it either reports a gap that does not exist or hides one. The same holds for a summary that silently leaves out part of the input, and for a figure given without its uncertainty.

Principle 10. Records are complete, protected and traceable.

Each run should leave a record that identifies its inputs, configuration, date and result, protected against later change, so that the result can be traced to its inputs and configuration and its integrity shown. Tracing does not mean producing the result again: a repeated run may differ, and a model version withdrawn by its provider cannot be run again. The ALCOA++ attributes apply.

Rationale.

A result that cannot be traced to its inputs and configuration cannot be defended in an audit, a dispute or a due-diligence review, whatever its quality.

Principle 11. Information about results and capabilities is clear and true.

Statements about a service, in reports, descriptions and conversations, should say exactly what was done, by whom, and with what limits. They should not claim more than has been measured.

Rationale.

Overstatement is a common quality failure in AI services and hard to repair once a user relies on it. Precise wording is also the strongest position, because it can be checked.

Principle 12. The practice learns and stays current.

Errors should be investigated to their root cause and corrected through a documented corrective and preventive action process. The regulatory and scientific sources that an AI tool or service relies on should be checked for currency before they are cited, and the right expertise should be involved for each domain.

Rationale.

Guidelines are revised, drafts become final, and AI tools fail in new ways. A practice that does not record and act on its own errors repeats them.

Alignment with the EMA and FDA guiding principles

EMA and FDA guiding principle (January 2026)	Where this guidance addresses it (in part where marked)
1 Human-centric by design	Principles 1, 4
2 Risk-based approach	Principle 3; 5.1, 5.5.6
3 Adherence to standards	Principles 4, 12; 5.2.9; section 6 (in part)
4 Clear context of use	Principle 2; 5.1
5 Multidisciplinary expertise	Principle 12; section 4; 5.11
6 Data governance and documentation	Principles 4, 10; 5.2, 5.8
7 Model design and development practices	Principles 6, 7; 5.4, 5.5, 5.6.6 (in part: model development itself is out of scope)
8 Risk-based performance assessment	Principle 7; 5.5, 5.7.6
9 Life cycle management	Principles 6, 7, 12; 5.4, 5.5, 5.12
10 Clear, essential information	Principles 8, 9, 11; 5.6, 5.9 (in part)

3 Glossary

Term	Meaning in this guidance
Advisory output	Output that informs a qualified individual and is not itself an approval, a review record or a decision. Advisory output that a person adopts into a record becomes part of that record.
Build	A fixed, identifiable configuration of an AI service, including its software, model version, instructions, settings and the way any knowledge base or document collection is built and searched, that produced a given result. It is wider than a software build: a change to instructions alone gives a new build.
Change control	The documented process by which a change to a build is assessed, tested and approved before use.
Computerised system	A system used in a clinical trial to collect, generate, process, hold or report data or records. This guidance treats an AI service used in a trial as one. What is expected of computerised systems is set out in ICH E6(R3), Annex 1, section 4.3, which does not address AI specifically.
Context of use	The specific role and scope of an AI service in a given task, including who relies on its output and for what decision.
Entrusted data	Data and documents of another party, or of trial participants, that a service processes.
Intake check	The checks applied to an input before any processing, to confirm the service can handle it safely.
Intended use	The written statement of what a service is for, what it takes as input and produces, and who uses its output.
Not assessable	A status meaning the tool could not reach a reliable result on an item, for example because the input could not be read or was incomplete, or the task lay outside what the service has been shown to do. It is a statement about the tool, not about the material under review.
Not met	For services that assess material against requirements: a status meaning the material under review (a document, a data set, a process description) does not meet a requirement or recommendation, or contradicts itself. It is a statement about that material.
Operator	Whoever builds or operates an AI service, whether the responsible party itself or a service provider. The term is not used in the sense of the EU AI Act.
Qualified individual	A person qualified by education, training and experience to perform the task in question.
Reference material	A fixed set of documents or data, not containing entrusted data, on which builds are tested and compared.

Term	Meaning in this guidance
Regression test	A test that runs a changed build on reference material and compares its results with the expected results, where these exist, and with the previous build, against criteria set in advance that take run-to-run variation into account.
Release	Of a build: the decision that a build may be used in production, taken because it has passed the defined conditions. Of output: the act by which a named role approves a result for use.
Responsible party	As in ICH E6(R3): the sponsor or the investigator, who remains responsible for the trial activity a service supports.
Run record	The protected record of a single run (one execution on an input, or one session of use): inputs, build, date, outputs and integrity data.
Service provider	As in ICH E6(R3): a person or organisation, such as a CRO or a technology company, that provides a service used by the sponsor or investigator.
Statement	A single assertion, figure, summary, recommendation or draft passage in an output on which a reader may rely.
Tolerance band	The range of run-to-run variation accepted for a type of result before it is investigated. It is similar in purpose to the pre-specified acceptable ranges of ICH E6(R3) (Annex 1, section 3.10.1.3), but applies to a service rather than to a trial.
User	The person or function that relies on the output of a service, for example a medical writer, a monitor, a data manager or a quality lead.
Variation (run-to-run)	Differences in output when the same build processes the same input more than once.
Withheld statement	A statement the service produced that was not released because it could not be supported by evidence or did not pass review.

4 Roles and responsibilities

Roles and responsibilities should be clear and documented before a service is used, as ICH E6(R3) expects for clinical trials (ICH E6(R3), Principle 10). One person may hold more than one role in a small organisation, provided the resulting conflicts are recognised, disclosed to those who rely on the output and controlled (see 4.3). Where a responsible party builds or operates an AI service itself, it holds both roles. Terms are defined in section 3.

4.1 Responsible party and operator

The responsible party should:

- define and document the intended use of each service and confirm that it matches how the output will be used;
- assess, in proportion to risk, whether the service is fit for its intended use, including a service built by others (see 5.5.6 and 5.5.7);
- decide whether output will become part of a trial record and, where it does, apply its own requirements for computerised systems and essential records;
- maintain oversight of any service provider involved (see 5.10);
- keep each decision taken on the output with a qualified individual.

The operator should:

- control the builds it uses and release only builds that have passed the tests set for them;
- handle entrusted data and documents under the agreed confidentiality and data protection terms;
- have output checked before it is relied on and, where the intended use calls for expert review, have named experts do so;
- keep run records and make the information needed to understand a result available to the responsible party;
- state the limits of each service and each result;
- investigate errors and complaints, correct their causes, and report incidents that may affect participants, trial data or confidentiality to the responsible party without delay.

4.2 Management

The management of the responsible party, and of any service provider for its part, is accountable for the service and for approving changes to it. Significant decisions about the service should be recorded with the question, the alternatives considered, the decision, its reasons and who took it.

4.3 Release reviewer

The release reviewer is the role that checks output before it is relied on outside the team that produced it. The reviewer confirms that each released statement carries its evidence, that withheld statements and items not assessable are reported, and that the output contains no claim beyond what was done.

In a small organisation, the person responsible for the service may also be its release reviewer, with no independent second reviewer. Where this is the case, it should be mitigated, for example as follows:

- every statement is tied to evidence that the user can check;
- statements withheld before release are disclosed by count;
- the user can reject any statement;
- where the intended use calls for expert review, named experts assess the statements in their field;
- the release review is documented, sampled under defined rules, and errors are handled through corrective and preventive action.

An independent second review should be introduced when the organisation's size allows.

4.4 Subject-matter experts

Where the intended use calls for expert review, named experts with relevant qualifications assess the statements in their field. Their names and roles should be stated in the output or in the record of its use. Experts should work under written confidentiality terms and follow the responsible party's data handling rules, including not entering entrusted data into other AI tools.

4.5 User

☐ Use a tool that others built

Responsibility for the trial stays with the responsible party. The user should:

- confirm that the intended use matches how it will use the output;
- check, under the quality system that applies, whether a planned use brings the output into a trial record, and make sure the responsible party's requirements are then applied;
- enter or provide only the data the intended use allows, through a route agreed in advance;
- use only services approved for the purpose (see 5.1.7), not private accounts or other tools;
- have qualified individuals review the output before acting on it.

4.6 Subcontractors and third parties

Any third party that processes entrusted data on the responsible party's behalf, including providers of computing and AI infrastructure, acts under a written agreement and, where personal data are processed, as a processor or sub-processor as the case may be. Whoever engages a third party remains responsible for that part of the service (see 5.10).

5 Lifecycle requirements

☐ Use a tool that others built

This section turns the principles into requirements for each stage of a service's life. Each subsection names the principles it implements. "Requirement" is used here in the sense of 1.4: what the author recommends, written with "should". None of it is binding unless a law, a guideline or a contract makes it so.

Who applies a requirement depends on the role. The operator applies the requirements directly. A responsible party that uses a service built by others should obtain evidence of the relevant controls from the service provider, or assess the service itself in proportion to risk, and record the result. The table shows the usual allocation; it is illustrative.

The depth of each requirement follows risk (Principle 3). A small organisation that uses a bought tool for lower-risk work can meet many of the requirements with a short written rule, the provider's evidence and a record of what was checked. An organisation with established processes for computerised systems, supplier qualification, change control, training and CAPA can meet the requirements through those processes; a separate system is not needed.

Subsection	Operator	Uses a service built by others
5.1 Intended use and risk	Applies	Applies
5.2 Confidentiality, data protection and access	Applies	Applies: choice of service, data allowed into it, access
5.3 Input qualification	Applies	Evidence from the provider; own rules on permitted inputs
5.4 Configuration and change control	Applies	Evidence from the provider; notice of changes (5.4.4)
5.5 Performance, fitness for purpose and monitoring	Applies to its part	Applies: fitness for purpose rests with the responsible party (5.5.7)
5.6 Output integrity and evidence	Applies	Applies as far as the service allows; gaps recorded
5.7 Human oversight, release and systems that act	Applies	Applies
5.8 Records, traceability and ALCOA++	Applies	Applies to own records; provider records by agreement
5.9 Communication of results and claims	Applies	Applies
5.10 Service providers and oversight	Applies to its own suppliers	Applies
5.11 Competence and training	Applies	Applies
5.12 CAPA and continual improvement	Applies	Applies

5.1 Intended use and risk (Principles 2, 3)

5.1.1 Each service should have a single written intended use. The same text should be used in the contract or internal specification, in the service description and, where the service produces a formal output, at the head of that output, so that no two versions diverge.

5.1.2 The intended use should state what the service takes as input, what it produces, who is expected to use the output, and that the output is advisory where that is the case.

5.1.3 Requirements should be assigned by asking which of three layers they belong to: law that applies regardless of GxP status (for example data protection and AI law); GxP requirements that apply only if the output becomes a regulated record or carries a regulated decision; and quality measures the responsible party adopts because the user needs to trust the result. Whoever builds, provides or uses a service should determine and record its own role and the classification of the service under the AI law that applies.

5.1.4 Whoever supplies a service to another organisation should state its own position on that service. It should not decide on the user's behalf that the user's quality system does not apply. Risk-based frameworks such as those of EMA and FDA consider, among other factors, how much the output influences a decision and what that decision affects. Review by a qualified individual can lower the first factor; it does not by itself take a service out of their scope, which depends on the use.

5.1.5 Activities within a service should be classified by the consequence of an error, taking into account how likely an error is and how readily a reviewer would detect it. Output used only inside the team sits at the lower end; output that leaves the organisation, and actions that cannot be undone or that involve entrusted data, sit at the higher end. Control should increase with the class, and in doubt the higher class applies.

5.1.6 Where oversight of lower-risk activities is delegated or sampled, the sampling rate and the error threshold that triggers closer control should be defined in advance and the error rate measured.

5.1.7 The responsible party should keep a list of the AI services in use, including AI functions switched on inside existing systems, each with its intended use, risk class and owner, and should make known which services are approved for which uses. Approval is given for a use, not for a tool. Using an approved service for a task of a higher class, for example drafting or reviewing a protocol or a medical rationale with a general office assistant, is a new use: it should have its own intended use and assessment before it starts.

5.1.8 Where a service supports a trial activity, its use should be considered in the risk management of that trial (see ICH E6(R3), Annex 1, section 3.10.1.1).

5.2 Confidentiality, data protection and access (Principle 4)

5.2.1 Where data pass between organisations, a confidentiality agreement should be in place before any entrusted data or document is received. Where personal data may be present, a data processing agreement should also be in place, naming the sub-processors. Whether a data protection impact assessment or safeguards for transfers to third countries are needed should be checked before use. Inside one organisation, the internal access and data protection rules take their place.

5.2.2 Where entrusted data and documents pass between parties, they should travel by a route agreed in advance that protects their confidentiality and integrity, such as a data room or an encrypted transfer.

5.2.3 Entrusted data and documents should be stored in a defined location with encryption, controlled access and known locations of storage, processing and support access.

5.2.4 Personal data in entrusted material should be minimised before any processing that leaves the environment controlled by the party holding the material. The effectiveness of the measures used for this should be measured, not assumed.

5.2.5 Each class of entrusted data should have a defined retention period and deletion trigger. Copies outside the primary store, such as backups, provider-side logs, caches and conversation histories, should be known and should expire within a stated period.

5.2.6 Work to improve a service should use public or synthetic reference material. Where a problem appears on entrusted material, it should be reproduced on reference material where possible; where that is not possible, work on the entrusted material stays under the terms of 5.2. Internal analysis of production runs should use aggregates, identifiers and locations rather than entrusted content, wherever this is possible.

5.2.7 Access to the service and to entrusted data should be based on roles, granted and withdrawn under a documented process, and reviewed periodically (see ICH E6(R3), Annex 1, section 4.3.8). Accounts should be individual. Administrative rights and technical credentials such as API keys should be limited to those who need them, stored securely and replaced when people or roles change. Access by the provider's staff should be covered by the agreement.

5.2.8 Where a service may process information that could reveal treatment allocation, the intended use should state which roles may receive its output. The risk of unblinding should be part of the risk assessment, and access that could unblind should be monitored so that any unblinding, intended or not, is detected and recorded (see ICH E6(R3), Annex 1, section 4.1).

5.2.9 Security measures should be proportionate to risk and cover threats specific to AI, such as instructions hidden in input documents, disclosure of data through connected tools, access by one user to another user's data, and weaknesses in components taken from third parties, such as models, libraries and connectors. Measures should rest on design, such as limiting what the service can reach, as well as on testing. An operator should test them before a build is released and after relevant changes; a party that uses a service built by others should obtain the provider's evidence.

5.3 Input qualification (Principle 5)

5.3.1 Each input should be checked at intake, before it is processed, for properties the service depends on: for example, that the input can be read reliably and completely, that it is within the size the service has been tested on, and that the confidentiality safeguards can operate on it.

5.3.2 Where a safeguard has not been shown to work on a class of input, that class should be refused at intake. The refusal should be recorded with a reason the user can understand and act on, such as supplying the input in a different format. A refusal rule should be lifted only on new measurement showing that the safeguard works on that class, recorded as a decision.

5.3.3 Intake checks should themselves be tested with positive and negative controls before they are relied on.

5.3.4 Before material is entered into a service, whoever supplies it should confirm that it may be used in this way, including the licence terms and copyright of third-party content such as dictionaries, rating scales or publications.

5.4 Configuration and change control (Principle 6)

5.4.1 Each build used in production should be identified by a unique reference recorded in every run record. For a third-party model, the reference should include the most precise version identifier the provider exposes; where none is exposed or fixed, the run record should say so.

5.4.2 A change that can alter results, including a change to instructions given to an AI model, should be treated as a change to the result and pass regression testing on reference material before production use.

5.4.3 Development builds, demonstration builds and builds released for production use should be distinct states. A build should move to the released state only because it has passed the defined release conditions. A released build later found to have a defect that affects its fitness should be withdrawn from production use, labelled as such, and its identification and its results on reference material kept for comparison.

5.4.4 Changes made by third parties, such as model versions, default settings, newly activated functions and sub-processors, should be assessed in the same way as internal changes and tested where they can alter results. Where a provider offers fixed model versions, the build should name one. Where a model version cannot be fixed, this should be stated to those who rely on the output. Where a change takes effect before it can be tested, it should be assessed afterwards, and monitoring under 5.5.5 should be designed to detect its effect. Announced end dates of model versions should be planned for as changes.

5.4.5 Where a service retrieves from a knowledge base or document collection, sources should be admitted under defined rules, checked for currency and versioned, so that each run record identifies the state of the collection used. A change to how the collection is built or searched is a change to the build; a change to its content under the admission rules need not be. A user should receive through the service only what that user is meant to access; existing access rights to connected sources should be reviewed before a service is connected to them. Deletion under 5.2.5 should include derived forms such as indexes.

5.4.6 When a service is retired, the responsible party and any service provider should agree in advance how run records and other records needed for the retention period stay available and readable, how entrusted data are returned or deleted, and how access is removed.

5.4.7 Where a trial activity depends on a service, the responsible party should have a contingency for the case that the service fails or is withdrawn, so that data essential to participant safety, trial decisions or trial outcomes are not lost or out of reach (see ICH E6(R3), Annex 1, section 4.3.6).

5.5 Performance, fitness for purpose and monitoring (Principles 3, 7)

5.5.1 Performance should be measured on a defined set of reference material covering the input types the service accepts, against criteria set in advance. Where the output concerns people, it should also be measured whether performance differs between relevant groups, on reference material that reflects the population concerned. A responsible party may also test a service on its own material in its own environment.

5.5.2 Run-to-run variation should be measured by repeated runs of the same build on the same inputs. The unit of comparison should be stated whenever a figure is reported.

5.5.3 Each type of result should have a tolerance band, set on reference material. A build that exceeds it in testing or monitoring should not be used in production until a person has examined the cause. Where the intended use calls for it, an input may be run more than once and differing items examined before the output is released. Items known to vary more than others should be marked as such in the output rather than silently removed or downgraded.

5.5.4 Figures describing performance should be stated with their basis. A figure computed on a set selected for the property being measured should not be presented as an out-of-sample result. Material used to develop or adjust a service should be kept apart from the material on which its performance is reported; where it is not, the figure should say so. Reference material should be versioned and each figure should name the version. A figure measured on published material may not hold for unpublished material and should be labelled accordingly.

5.5.5 Monitoring should continue after a build is released, with each new build and at defined intervals.

5.5.6 This guidance treats an AI service used in a clinical trial as a computerised system in the sense of ICH E6(R3). Its fitness for purpose should be established and kept by a risk-based approach that considers its intended use, the importance of the data or records it affects, and its potential to affect participants and the reliability of trial results (see ICH E6(R3), Annex 1, section 4.3.4). The controls in 5.3 to 5.5 can form part of the evidence for this.

5.5.7 The responsible party is responsible for the validation status of a system it uses throughout its life cycle and should ensure that it is validated as fit for purpose for its use in the trial, including a system developed by others (see ICH E6(R3), Annex 1, section 4.3.4). Where output becomes part of a trial record, the responsible party's approach to computerised systems and records applies.

5.6 Output integrity and evidence (Principles 8, 9)

5.6.1 Each statement should point to the input it rests on, precisely enough to be checked, and, where a standard, guideline or internal consistency rule is applied, name it. For a document, this means the location (for example page and section) and a verbatim quote, confirmed against the input by a means independent of the model that produced it; for data, the records used and the calculation; for a generated text, the sources used.

5.6.2 A statement that cannot be supported in this way should be withheld or clearly marked as unsupported. A quote that cannot be found in the input makes the statement unsupported. Where the service selects what it shows, outputs should state how many statements were withheld and for which reason (unsupported, or not passed in review).

5.6.3 A service that assesses material against requirements should define in advance how a missing or conflicting element is graded and which source each grading relies on, and should use the statuses defined in section 3 consistently.

5.6.4 Statements should be worded as what was or was not identified in the input as provided, not as facts beyond it, and should cite guidelines as recommending or requiring according to their own wording. Draft guidelines may be cited as an indication, never as a requirement.

5.6.5 Items the tool could not assess should be reported as not assessable, with the reason, including where a processing step failed or was incomplete, and kept apart from negative results about the material or data concerned.

5.6.6 The coverage of each output should be visible: which parts of the input were supplied to the model and which were left out, shortened or unreadable. Where the way the service reached a result cannot be inspected, the output should say so, and the evidence under 5.6.1 should allow the result to be checked independently of the service's reasoning.

5.6.7 Output should show which parts are produced with AI assistance, which are taken unchanged from a source, and who reviewed it. Labelling should be specific rather than a blanket statement.

5.7 Human oversight, release and systems that act (Principle 1)

5.7.1 Output should be checked by a qualified individual before it is relied on, to a depth that follows its risk class (see 5.1.5 and 5.1.6); where the intended use calls for expert review, by named experts with relevant qualifications. The output, or the record of its use, should state who checked what.

5.7.2 Where output is reviewed statement by statement, the reviewer's own status for each statement (accepted, changed, rejected, not reviewed) should be recorded separately from the tool's result.

5.7.3 The responsible party should define who may release output of each risk class for use outside the team that produced it.

5.7.4 Numbers passed on to users or third parties should come from a calculation that can be repeated, not from the wording of a language model alone, or be clearly labelled as checked or computed by a person, with method and inputs.

5.7.5 Where the person responsible for a service is also its release reviewer (see 4.3), this should be stated to those who rely on the output, and the controls used to address the missing independent review should be named.

5.7.6 The effectiveness of human review should be checked from time to time, for example by measuring errors found after release and the share of statements reviewers change or reject.

5.7.7 Where a service can act on other systems (for example write, send, change or delete), each permitted action should be listed in the intended use and limited by technical permissions to what is listed. Actions should be logged as in 5.8. The service should act under an identity that has no more rights than the listed actions need. Content that a service reads is data, not instruction; as testing alone cannot assure this for language models, a service that reads content from outside should not be able to send entrusted data out without confirmation by a person.

5.7.8 Actions of the higher classes in 5.1.5 should take effect only after confirmation by a person.

5.8 Records, traceability and ALCOA++ (Principle 10)

5.8.1 Each run should produce a record that identifies the input (for example by cryptographic hash), the build, date and time, and the outputs. Where the input is deleted under 5.2.5, the record should state who keeps the input needed to trace the result, normally the responsible party. Where a run includes retrieval or actions on other systems, the record should identify what was retrieved, in which version or at what time, and what was done and when.

5.8.2 Completed run records should be protected against change. A correction should create a new record that references the original, never overwrite it. Protection against change does not prevent deletion under 5.2.5; a deletion should itself be recorded.

5.8.3 Integrity should be demonstrable, for example by technical means that reveal any later change and that those who administer the records cannot alter unnoticed.

5.8.4 Actions on output that affect a trial record (adoption, change, rejection) should be captured with who, when and, where relevant, why, in line with the audit trail expectations of ICH E6(R3), Annex 1, section 4.2.2, and should be attributable to an individual (section 4.3.8).

5.8.5 Records should be attributable, legible, contemporaneous, original and accurate, and also complete, consistent, enduring, available when needed and traceable (ALCOA++). Reconstructed entries should be labelled as reconstructions.

5.8.6 Where a bought service does not produce a run record as described in 5.8.1, the record may consist of the provider's logs and the user's own note of what was used (service, date, input, output adopted); the gap should be recorded. Run records that contain entrusted content fall under 5.2. Their retention period should be defined; it need not equal that of the trial record they support. Provider logs needed beyond the provider's retention should be exported.

5.9 Communication of results and claims (Principle 11)

5.9.1 Every factual statement in an external document about a service or its results should be checked against its primary source before release.

5.9.2 Those who supply or describe a service should not attribute a formal regulatory or certification status to a service unless that status has been formally established, and should not promise identical results on every run.

5.9.3 Statements about data handling should describe only what is true and in place at the time of publication.

5.9.4 Where output of a service may be relevant to a regulatory submission or decision, the responsible party should consider what the authority concerned expects to be told and whether to seek its advice early. The EMA reflection paper listed in section 6 advises early interaction where an AI system may affect the benefit-risk balance of a medicinal product.

5.10 Service providers and oversight (Principles 4, 10)

5.10.1 Before a service provider is used, its suitability should be assessed. The responsible party should have access to the information needed for oversight, for example procedures, performance measurements and notice of changes, and should agree in writing on incident reporting, further subcontracting, access for audits and inspections or, where the provider does not grant it, independent assurance reports, the location of processing, and what applies when the service is unavailable, changed or ended, including how data and records are handed back (see ICH E6(R3), Annex 1, section 3.6). Third parties that process personal data should be listed as sub-processors.

5.10.2 Data protection settings that a responsible party relies on, such as data retention, training use and processing region, should be verified by evidence rather than taken from a description: the contract terms, the configuration as shown in the administration console, and independent assurance reports. The evidence should state which products, functions and interfaces it covers, be kept, and be checked again after relevant changes. What cannot be verified technically should be recorded as relied on by contract.

5.10.3 The terms of third-party services and licences should be followed by wording and purpose.

5.10.4 Where a provider cannot give an assurance or the evidence for it, the responsible party should record the gap and decide in proportion to risk: limit the use or the data allowed into the service, add controls of its own, accept the remaining risk by a documented decision of a named person, or choose another service. An assurance that the law requires, such as a data processing agreement where personal data are processed, cannot be replaced in this way.

5.11 Competence and training (Principles 1, 12)

5.11.1 People who build, operate or review an AI service should have documented training appropriate to their role, including AI literacy (see also Article 4 of the EU AI Act as amended in 2026).

5.11.2 Expertise outside the responsible party's own field, such as biostatistics, a therapeutic area or law, should be obtained from qualified people and their role stated.

5.11.3 Attendance at training should not be presented as proof of competence; where competence matters, it should be assessed.

5.12 CAPA and continual improvement (Principle 12)

5.12.1 Errors, near misses and complaints should be recorded, investigated to their root cause and addressed by corrective and preventive actions with owners and dates. Corrections to earlier statements, including one's own, should be recorded openly. Incidents that may affect participants, trial data or confidentiality should be reported to the responsible party without delay, and in time for the responsible party to meet its own reporting obligations, for example under data protection or clinical trial law.

5.12.2 Lessons that apply beyond the case should become rules, and rule changes should reach every person and process that applies them.

5.12.3 Regulatory and scientific sources should be checked for currency before each citation: final or draft status, latest revision, and effective date. Statements in regulatory assessments should distinguish what a source states from what is inferred or assumed.

6 References

Status checked on 24 November 2026. The list contains documents cited in this guidance and further guidance relevant to its subject. Draft documents are cited for awareness only and are not treated as requirements.

Document	Issuer, reference	Version and date	Status
Guideline for Good Clinical Practice E6(R3)	ICH; EU: EMA/CHMP/ICH/135/1995 (Step 5); consolidated with Annex 2: EMA/CHMP/ICH/162879/2026	Principles and Annex 1: ICH Step 4, 6 January 2025. Annex 2: ICH Step 4, 3 June 2026	Final. In effect in the EU since 23 July 2025 (Principles, Annex 1); Annex 2 from 15 January 2027
Guideline on computerised systems and electronic data in clinical trials	EMA, GCP Inspectors Working Group; EMA/INS/GCP/112288/2023	Final version adopted 7 March 2023	Final. Published 9 March 2023; in effect six months after publication (September 2023)
General Considerations for Clinical Studies E8(R1)	ICH; EU: EMA/CHMP/ICH/544570/1998	Step 4, 6 October 2021	Final. In effect in the EU since 14 April 2022
Addendum on Estimands and Sensitivity Analysis E9(R1)	ICH; EU: EMA/CHMP/ICH/436221/2017	Step 4, 20 November 2019	Final. In effect in the EU since 30 July 2020
Quality Risk Management Q9(R1)	ICH	Step 4, 18 January 2023; EU: EMA/CHMP/ICH/24235/2006, in effect since 26 July 2023	Final
General Principles for Model-Informed Drug Development M15	ICH; EU: EMA/CHMP/ICH/496426/2024	Step 4, 29 January 2026	Final. In effect in the EU since 23 July 2026
Clinical Electronic Structured Harmonised Protocol (CeSHarP) M11	ICH; EU: EMA/CHMP/ICH/778799/2022	Step 4, 19 November 2025	Final. In effect in the EU since 11 June 2026
Guiding Principles of Good AI Practice in Drug Development	EMA and FDA, joint	January 2026 (published 14 January 2026)	Published principles; not guidance; non-binding
Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle	EMA (CHMP, CVMP); EMA/CHMP/CVMP/83833/2023	Adopted by CHMP on 9 September 2024 and by CVMP on 11 September 2024	Final reflection paper
Considerations for the Use of AI to Support Regulatory Decision-Making for Drug and Biological Products	FDA (CDER, CBER and others), docket FDA-2024-D-4689	January 2025	Draft guidance; not for implementation

Document	Issuer, reference	Version and date	Status
Regulation (EU) 2024/1689 (Artificial Intelligence Act)	European Parliament and Council	Regulation of 13 June 2024; consolidated version of 27 July 2026	In force. Article 4 (AI literacy) as replaced by Regulation (EU) 2026/1744
Regulation (EU) 2026/1744 (Digital Omnibus on AI), amending Regulation (EU) 2024/1689	European Parliament and Council; OJ L, 24.7.2026	Regulation of 8 July 2026	In force since 27 July 2026. Moves the application of the high-risk requirements to 2 December 2027 (Article 6(2), Annex III) and 2 August 2028 (Article 6(1), Annex I)
Regulation (EU) 2016/679 (General Data Protection Regulation)	European Parliament and Council	Regulation of 27 April 2016	In force since 24 May 2016; applies since 25 May 2018
EudraLex Volume 4, Annex 11 Computerised Systems	European Commission	2011 version	In force. Revised draft consulted 7 July to 7 October 2025; no final version published
EudraLex Volume 4, Annex 22 Artificial Intelligence	European Commission	New draft for consultation, 7 July to 7 October 2025	Draft; no final version published. GMP scope
GAMP 5 Guide, Second Edition	ISPE	Second Edition, July 2022	Current
GAMP Guide: Artificial Intelligence	ISPE	Published July 2025	Current
GXP Data Integrity Guidance and Definitions	MHRA	Revision 1, March 2018	Current
Good Practices for Data Management and Integrity in Regulated GMP/GDP Environments	PIC/S, PI 041-1	Adopted 1 June 2021, in force 1 July 2021	Current

☐ Use a tool that others built

Annex A: Examples of use and where their risk tends to lie (informative)

This annex is not part of the standard.

The examples are illustrative. They do not classify any use. The class of a use depends on the trial, on what the output influences and on how it is checked. The responsible party decides it (see 5.1.5), starting from the factors critical to the quality of the trial (ICH E6(R3), Principles 6 and 7; ICH E8(R1)).

Example of use	What an error could affect	Tends towards
Literature search and summaries for orientation inside the team	The team's own understanding; decisions are taken on the sources	Lower end
Meeting notes and wording of internal drafts	Internal communication	Lower end
Advisory review of a protocol or plan for gaps and inconsistencies	Trial design, if a finding is missed or a wrong one is adopted	Middle
Country and site feasibility research	Planning and recruitment assumptions; which groups of patients are considered	Middle
Draft text for monitoring reports or data queries	Trial records, once the text is adopted	Middle
Support in assessing eligibility of participants	Participant safety and rights; who is included	Higher end
Handling of safety information	Participant safety; reporting obligations	Higher end
Coding, cleaning or derivation of data for an endpoint	Reliability of trial results	Higher end
Text or analyses for a regulatory submission	Regulatory decisions	Higher end
Any use that could reveal treatment allocation	Blinding (see 5.2.8)	Higher end

Build, operate or assess

Annex B: Where ICH E6(R3) addresses related topics (informative)

This annex is not part of the standard. It helps readers who know ICH E6(R3) to find their way. It lists where E6(R3) addresses a related topic for clinical trials in general. It is not a statement of conformity or of equivalence. Section numbers refer to the Principles and to Annex 1 of ICH E6(R3) in the version listed in section 6.

This guidance	ICH E6(R3)
4 Roles and responsibilities	Principle 10; Annex 1: 3.3 Allocation of activities, 3.6 Agreements
5.1 Intended use and risk	Principles 6 and 7; Annex 1: 3.10 Quality management, 3.10.1 Risk management, 4.3 Computerised systems
5.2 Confidentiality, data protection and access	Annex 1: 2.12 Records (investigator), 4.1 Safeguard blinding in data governance, 4.2.5 Data transfer, exchange and migration, 4.2.8 Destruction, 4.3.3 Security, 4.3.8 User management
5.3 Input qualification	No direct counterpart
5.4 Configuration and change control	Annex 1: 4.3.4 Validation, 4.3.6 System failure
5.5 Performance, fitness for purpose and monitoring	Annex 1: 3.10.1 Risk management, 4.3.4 Validation
5.6 Output integrity and evidence	Principle 9; Annex 1: 4.2 Data life cycle elements
5.7 Human oversight, release and systems that act	Principles 5 and 10; Annex 1: 3.9 Sponsor oversight, 3.11 Quality assurance and quality control
5.8 Records, traceability and ALCOA++	Annex 1: 3.16.3 Record keeping and retention, 4.2.2 Relevant metadata, including audit trails, 4.2.7 Retention and access
5.9 Communication of results and claims	No direct counterpart
5.10 Service providers and oversight	Annex 1: 3.6 Agreements
5.11 Competence and training	Principle 5; Annex 1: 3.4 Qualification and training, 4.3.2 Training
5.12 CAPA and continual improvement	Annex 1: 3.10.1 Risk management, 3.12 Noncompliance

Annex C: Questions by function (informative)

This annex is not part of the standard. It lists questions that the functions of a sponsor, a CRO or a site tend to ask when AI enters their work, and where this guidance answers them or gives a direction. In a small organisation one person may hold several of these functions; in a large one each has its own department. The answers are the same; the depth follows risk (Principle 3).

Function	Typical question	Where to look
Management	Where do we start, and how much is enough?	At a glance; start of section 5; 5.1.5; 5.1.7
Clinical operations	May the team use an assistant for monitoring reports or queries? What about the tools our CRO uses?	5.1; 5.7; 5.10; Annex A
Investigator and site	May site staff use AI on trial data or with participants?	4.1; 5.2; 1.2 (AI that interacts with participants is not addressed)
Data management and biostatistics	What applies when AI supports coding, cleaning, derivation or analysis?	Annex A; 5.5; 5.8; 1.3 (ICH M15)
Medical writing	What has to be checked and labelled in an AI-assisted document?	5.6; 5.7; 5.3.4
Pharmacovigilance	May AI support the handling of safety information?	Annex A; 5.1.5; 5.7; 5.12.1
Regulatory affairs	Which guidance applies, and what do we tell the authority?	1.3; 5.9.4; section 6
Quality assurance	Does the tool need validation? What do we show in an audit or inspection?	5.5.6; 5.5.7; 5.1.7; 5.8; Annex B
IT and information security	Which tools are approved? What happens when a model changes or a service fails?	5.1.7; 5.2.9; 5.4.4; 5.4.7; 5.10.2
Data protection and legal	Which agreements and assessments are needed? May we enter third-party content?	5.2; 5.1.3; 5.3.4; 5.10.2 to 5.10.4
Procurement and vendor management	What do we ask a provider, and what if it cannot answer?	5.10; table at the start of section 5
Training	Who needs which training, and what counts as competence?	5.11

Version history

Version	Date	What changed and why
1.0	26 November 2026	First edition

Imprint

© 2026 Dr. Christoph Franz. Published by Dr. Franz Clinical Consulting GmbH.

Dr. Franz Clinical Consulting GmbH, Johannes-Brahms-Str. 4a, 68542 Heddesheim, Germany. Registered at Amtsgericht Mannheim, HRB 760280. Managing director: Dr. Christoph Franz. Contact: info@drfranzclinical.com

How this document was made: the principles, the scope and all decisions on content are the author's. The text was drafted and revised with AI tools under the author's direction. The citations were checked against the primary texts with AI tools, including an independent AI review, and a sample was checked by the author. The author is trained in ICH E6(R3), has read the EMA and FDA guiding principles and reviewed the whole text. He holds editorial responsibility and has not read every other cited source in full.

How the author applies this standard: Dr. Franz Clinical Consulting GmbH (DrFC) uses this standard as its reference for the AI services it provides. At the date of this edition, DrFC's founder is both responsible for the services and their release reviewer; there is not yet an independent second reviewer inside DrFC. How DrFC addresses this is set out in each client contract. The implementation is documented internally and is not part of this document. This statement is not part of the standard and may change between editions.

Disclaimer

This document describes a practice standard. It is general information, not legal, regulatory or professional advice, and it does not replace an assessment of the individual case. Readers remain responsible for their own compliance with the laws and guidelines that apply to them. It has not been issued, reviewed or endorsed by ICH, EMA, FDA or any other authority.

Applying this document does not by itself demonstrate compliance with any law or guideline. The author and the publisher do not certify, audit or confirm conformity with it, and no one may state that they do.

The content was prepared with care. No warranty is given that it is complete, correct or current; the references in section 6 reflect the status on the date stated there.

The statement in the imprint on how the author applies this standard describes how the author works. It does not create rights for any reader. Obligations towards clients arise only from the contract concluded with that client.

To the extent permitted by law, the author and the publisher accept no liability for loss arising from the use of this document. This does not apply to intent or gross negligence, to injury to life, body or health, or where liability cannot be excluded by law.

DrFC Practice Standard

Good AI Practice in Clinical Research

DrFC

Dr. Franz Clinical Consulting GmbH · Heddesheim · Amtsgericht Mannheim HRB 760280 ·
Managing director: Dr. Christoph Franz
info@drfranzclinical.com