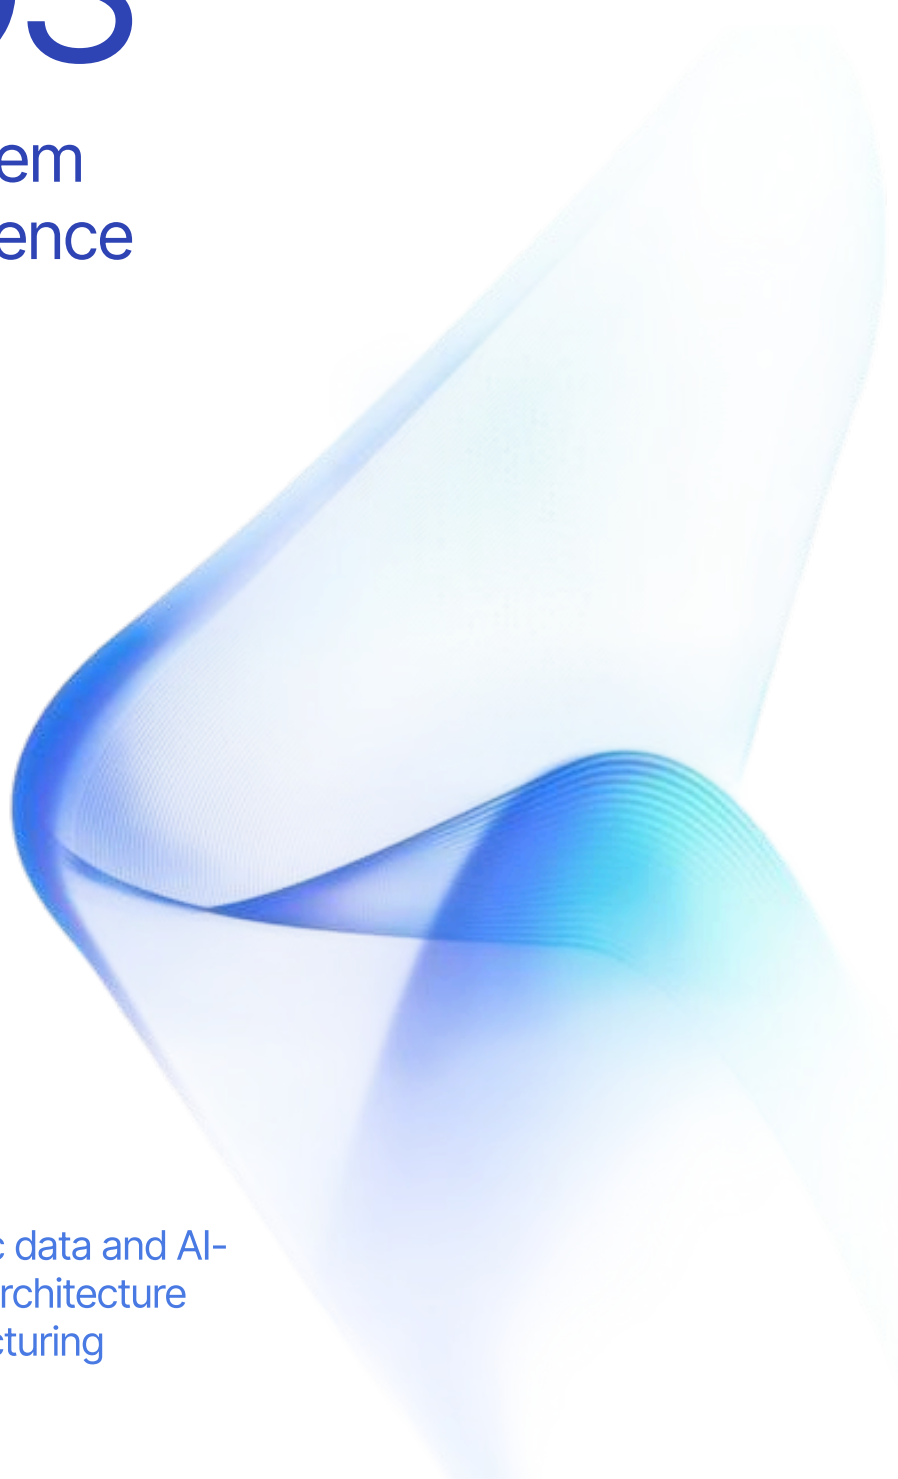


Tetra OS

The Operating System
for Scientific Intelligence

Industrializing AI-native scientific data and AI-enabled use cases — a unified architecture for biopharma R&D and manufacturing

A large, abstract graphic on the right side of the page, consisting of flowing, overlapping blue and white shapes that resemble a ribbon or a stylized wave, creating a sense of movement and depth.

THE PROBLEM

Science runs on fragmented data and a DIY mindset

Pharma's scientific data is born in captivity. Every instrument, ELN, and LIMS generates its own proprietary format — thousands of silos that should be the raw material of intelligence but are instead fragmented, opaque, and inert.

The result is Eroom's Law: drug development costs doubling every nine years even as compute doubles every eighteen months. Science runs on a pre-industrial operating system.

Closing that gap is not a tooling problem.

It requires an operating system purpose-built for science — one that converts raw instrument data into governed, reusable, AI-native assets and operationalizes them where science actually happens: at the bench, in the process, across the value chain.

10,000+

life sciences organizations worldwide

1M+

disconnected scientific data silos

500+

instrument & software vendors

\$2B · 10yr

to develop one new therapeutic

WHAT MAKES SCIENTIFIC DATA HARD

Five attributes every dataset must preserve

Successfully preparing data for AI applications requires more than moving files to the cloud.

It requires a thorough understanding of what scientific data actually is and what attributes must be preserved across its entire journey.

Five properties define whether scientific data is fit for AI — and each demands deep expertise to get right.



Complexity

A fragmented ecosystem of 500+ instruments, ELNs, informatics apps, robotics, and IoT sensors, each with distinct format and content requirements. Any work to prepare scientific data for AI must retain full information about instruments, methods, and processes.



Contextualization

Metadata gives scientific data meaning. Creating it requires subject-matter expertise across instruments, applications, scientific workflows, and use cases — not just data engineering.



Compliance

The life sciences industry is heavily regulated to ensure patient safety. Data must meet standards like US FDA 21 CFR Part 11 and GxP, with full traceability and integrity preserved at every step.



Control

All data belongs to the users who generate it — available to its originator to reuse and leverage, independently of vendor, format, status, or location, unconditionally.



Collaboration

Scientific data is generated across a globally distributed value chain of biopharmas, CROs, and CDMOs. Sharing it securely while preserving integrity and traceability is as complex as generating it — and equally critical.

An immutable order of operations

Raw data as generated by scientists and lab instruments can't be used for Scientific AI as-is. Value is created in a fixed sequence, an immutable order of operations that Tetra OS is built around. Each step raises the data's readiness for AI; skip one and intelligence resets with every project.



Raw Data

as generated

leaves the instrument, ELN, or LIMS: proprietary, unstructured, and locked to the vendor that produced it. It captures the experiment but carries none of the context a machine needs to reason about it.

In this state it is effectively inert: impossible to search, compare, or feed to a model at scale.



Replatformed

open and harmonized

The same data, collected and centralized in the cloud through validated, vendor-certified integrations. Provenance, audit trail, and regulatory context (21 CFR Part 11, GxP) travel with it, so nothing is lost in the move.

It is now liquid and accessible, and available to the people and systems that need it, rather than stranded on an instrument PC.



Engineered

open and harmonized

Replatformed data is lifted out of its proprietary container into open, vendor-agnostic formats, then harmonized against shared scientific schemas, taxonomies, and ontologies.

Results from different instruments and sites finally speak the same language and can be compared and combined — interoperable by any analytics or AI tool, without bespoke glue code.



AI-Native

governed and reusable

Harmonized data enriched with the scientific context that makes it reusable: linked to the right entities, use cases, and relationships, and governed so every consumer can trust it.

These are the building blocks Scientific AI actually runs on — the assets that feed agents, retrieval, and model training. Engineered once, they pay off across every future project.

Each layer compounds on the one before it.

Where the architecture is strong, every experiment makes the next one faster; where it is weak, intelligence resets with every project.

Tetra OS: four pillars, one operating system



Scientific Data Foundry

Converts proprietary, unstructured scientific data into AI-native schemas, taxonomies, and ontologies — creating governed, reusable scientific memory.



Scientific Use Case Factory

Productizes AI-enabled scientific workflows into validated, reusable assets that can be deployed across teams, sites, and domains.



Tetra Sciborgs

Embedded scientist-engineers who connect science, data, and IT to operationalize adoption and deliver measurable outcomes.



Tetra AI

The reasoning and orchestration layer — agents, RAG, and copilots that help scientists navigate complex workflows, surface context, and accelerate decisions.

Open, AI-native, and collaborative by design

Open & agnostic

Data transformed into open, vendor-agnostic JSON, accessible via REST API and SQL — usable with any best-of-breed application.

AI-native

Large-scale, liquid, compliant datasets with use-case taxonomies and ontologies, plus transparency into how AI reaches a conclusion.

Collaborative

Securely share contextualized data across biopharmas, CROs, CDMOs, and partners — fueling federated learning and model improvement.

Governed

Validated integrations preserve integrity, traceability, and compliance with 21 CFR Part 11 and GxP at every layer.

Built on a scientific lakehouse

Natively supports Snowflake, Databricks, and Amazon Athena & Redshift — enabling multi-modal consumption from discovery and analytics to high-performance computing and AI-model training.

PROVEN OUTCOMES

From replatformed data to measurable impact

When scientific data becomes AI-native, the results compound across R&D, manufacturing, and QA/QC.

UP TO

70%

reduction in clone selection time

UP TO

50%

increase in candidate success

MORE THAN

25x

faster data preparation

FROM

340

hours/month saved in chromatogram merging/overlay

What customers are saying



Our expanded partnership with TetraScience is delivering measurable value through unified access to instrument and CRO data that powers our automation and analytics at scale. The platform's audit capabilities have streamlined our regulatory preparation processes."



Linus Goerlitz,
Regulatory Science Transformation Lead



By transforming how our scientists access, analyze, and share research data, we're unlocking new levels of productivity and enabling AI-powered insights through a connected, online data environment. Beyond boosting productivity, we're leveraging data and agentic AI to accelerate innovation across our drug discovery engine."

Jim Villa,
Global Head of Research Strategy & Operations



Science is a compounding system — or it should be. Every experiment should make the next one faster and more accurate. Every workflow should get easier to replicate and improve. That is what Tetra OS is built to deliver: not a tool for the next project, but an architecture that makes every project better than the last.

See Tetra OS in your environment →