Pharma Growth Runs on Digital Platforms

How governance-first, trusted platforms turn IT into a regulated growth engine

Persistent
Re(AI)magining™ the World

Contents

New Mandate for Pharmaceutical IT

The pharmaceutical industry, an ecosystem historically defined by the meticulous pursuit of scientific breakthroughs and an uncompromising adherence to regulatory compliance, is currently navigating a fundamental and irreversible transformation in its relationship with Information Technology (IT).

For decades, IT was relegated to a defensive, often under-resourced, role — a necessary evil viewed primarily as a cost center dedicated to supporting utility, stability and compliance maintenance. This included rudimentary infrastructure management, data center operations and ensuring baseline adherence to foundational regulations such as GxP and HIPAA. That narrow, reactive model is now decisively obsolete.

The new paradigm elevates IT from a supportive function and compliance mandate to a strategically regulated growth engine. This is more than a title change; it is a profound shift in executive accountability. IT leadership is now directly responsible for delivering core business objectives: accelerating discovery, facilitating the deployment of personalized treatments and optimizing global operations — all while upholding the highest, most auditable standards of data integrity and regulatory adherence across a global footprint.

The contrast in executive mindset and strategic focus illustrates this changes.

Executives who struggle...

Are trapped in tactical, legacy thinking, asking infrastructure-focused questions: “Which cloud should we choose to save the most money?” or “How fast can we complete the physical migration of our servers?”

Executives who succeed...

Adopt a strategic, governance-first perspective, asking outcome-oriented questions: “What secure, validated platform do we need to design today that regulators, scientists and the business can explicitly trust for the next decade of discovery and commercialization?”

Trends in Cloud, Data and AI in Pharmaceutical Industry

AI value is poorly measured

60% of organizations fail to define or track financial KPIs for AI, making ROI unclear and slowing adoption (BCG).

AI remains stuck in pilots

Only 13% of pharma companies are piloting AI-enabled protocol design, with very limited scaling across the enterprise (Bain).

Cybersecurity is not embedded

61% of companies treat cybersecurity as a compliance checkbox rather than an integrated enterprise capability (KPMG).

Transformation will redefine competitiveness

Pharma and MedTech companies that fully embrace Cloud, Data and AI will operate fundamentally differently and more competitively within five years (McKinsey).

Hybrid cloud remains the norm

60% of pharma companies still operate across cloud and on-prem environments, driven by regulatory concerns and lack of strategic direction (PwC).

Data fragmentation is the biggest barrier

68% of pharma organizations cite data silos, inconsistent quality and security / privacy issues as major implementation challenges (KPMG).

Key Takeaways

- AI will not scale without secure cloud foundations, trusted data and measurable value
- A comprehensive Cloud, Data and AI strategy is now a business imperative — not an IT choice

The Dual Mandate — Regulation Meets Growth

The modern Pharma IT organization operates under a non-negotiable dual mandate, which forms the core tension of its strategy: maintaining (1) regulatory rigor while simultaneously driving; (2) business acceleration and innovation. Failure in either area results in catastrophic business impact, whether through scientific stagnation or regulatory penalties.

Regulatory requirements across global bodies (FDA, EMA, PMDA, etc.) are becoming exponentially more complex and granular, subjecting every IT system involved in the drug lifecycle, from early-stage research to post-market surveillance, to rigorous scrutiny.

The Regulated Core:

Regulatory requirements are not just foundational; they are becoming more complex and granular. IT systems now touch every critical, auditable stage of the drug lifecycle, making them subject to rigorous scrutiny by global bodies such as the FDA, EMA and regional authorities.

Regulatory Area	IT Impact and Requirement	Strategic Imperative
GxP Compliance	Computerized System Validation (CSV): IT must demonstrate that systems used in clinical trials (GCP), manufacturing (GMP) and lab processes (GLP) are fit for their intended use. This requires a robust System Development Life Cycle (SDLC), documented validation master plans and precise change control. The move to modern GxP methodologies like CSV-light and Agile GxP is critical for speed.	Quality Assurance of Processes
Data Integrity (DI)	ALCOA+ Principles: Absolute assurance that electronic records are Attributable, Legible, Contemporaneous, Original and Accurate (plus Complete, Consistent, Enduring, Available). DI failures are a major cause of FDA Warning Letters. IT must implement robust electronic audit trails, secure data handling procedures and manage data lifecycle from creation to archiving.	Trust in Scientific Evidence
Patient Data Privacy	Global Sovereignty and Consent: Strict adherence to GDPR, HIPAA and emerging regional laws (e.g., CCPA, Brazil's LGPD) mandates sophisticated data masking, anonymization techniques, access controls and transparent governance for cross-border data transfer, especially in decentralized clinical trials (DCTs).	Ethical and Legal Data Handling
Pharmacovigilance (PV)	Automated Safety Monitoring: Systems for collecting, processing and reporting Adverse Events (AEs) must be validated, highly reliable and capable of rapid scaling to manage global safety data. The transition to cloud-based PV systems is driven by the need for faster global reporting.	Patient Safety and Transparency

The Growth Engine: Driving Innovation and Speed to Market

The global competitive landscape demands that pharmaceutical firms shrink the timeline from discovery to commercialization. IT is no longer an afterthought; it is the primary accelerator of this innovation.

- Accelerated R&D and AI / ML: IT platforms must enable computational drug discovery. This requires managing and leveraging petabytes of heterogeneous data (genomics, proteomics, imaging) and operationalizing complex AI / ML models for target identification, lead optimization and even running in silico (computer-simulated) trials. This necessitates compliant, highly elastic High-Performance Computing (HPC) environments and industrial-scale data orchestration.
- DCTs and RWE: IT builds the secure, validated infrastructure that integrates data from patient mobile apps, wearable devices and remote monitoring tools with Real-World Evidence (RWE) streams (e.g., Electronic Health Records or pharmacy data). This platform approach enhances trial efficiency, broadens patient reach and generates richer, longitudinal patient insights.
- Personalized Medicine and Digital Health: Supporting targeted therapies requires IT to build compliant, federated data lakes that securely integrate patient genomic, proteomic and clinical data. Furthermore, IT must support the lifecycle and launch of Software as a Medical Device (SaMD) products (such as companion diagnostics or digital therapeutics), introducing entirely new regulatory burdens (such as IEC 62304 for medical device software).
- Intelligent Manufacturing and Supply Chain: Implementing IoT sensors, AI-driven predictive maintenance and blockchain solutions within GMP-regulated manufacturing environments ensures “vein-to-vein” or “plant-to-patient” tracking. This guarantees product quality, enables anti-counterfeiting measures and establishes a regulatory chain of custody that is auditable at every step.

Cloud is Not the Strategy — Control Is

A common mistake in pharma is treating cloud adoption as a tactical exercise focused on mere migration speed or cost reduction. In a regulated industry, this approach is incomplete and inherently dangerous. Cloud is fundamentally an accelerator and acceleration without commensurate control amplifies risk to an existential level.

Strategy is defined by who controls outcomes, the consistency of those controls and their absolute defensibility under regulatory and legal scrutiny.

That is why the cloud is not the strategy. **Control is**.

Regulation has Shifted from Documentation to Demonstrable Control

Regulators are moving past the static review of paper-based validation documents. They now demand that organizations demonstrate control in live operation and prove the continuous validity of their systems. This transforms regulatory scrutiny from a policy issue into a core architectural and engineering question:

- Can you show who accessed this data and why?
- Can you prove the system behaved consistently over time?
- Can you demonstrate how failures are detected and contained?
- Can you reproduce this result exactly?

These are not policy questions. They are architectural questions.

Only robust IT platforms, specifically those built on Infrastructure as Code and automation, can enforce controls automatically, generate immutable audit trails and reconstruct historical states on demand. This makes IT the primary interface and evidence provider for regulators.

Data: Core Asset and Existential Liability

Pharma's growth is entirely data-driven, relying on RWE, advanced analytics and AI / ML models. Concurrently, this data presents a massive, compounding liability due to decades-long retention obligations (sometimes >25 years), existential privacy violation risks and the need for rigorous lineage tracking to maintain regulatory confidence in scientific findings.

This duality — data as asset and liability — means that growth depends on governance. IT platforms determine whether data can be truly trusted, safely reused for subsequent research, deployed to support AI / ML and ultimately withstand any legal or regulatory challenge.

IT platforms determine whether data can be:

- Safely reused
- Used to support AI without introducing risk
- Leveraged to withstand legal or regulatory challenge
- Trusted by downstream systems

In other words, data does not create growth by existing. It creates growth only when IT platforms can prove data integrity, purpose and control.

AI Turns IT from Enabler into Decision-Shaping Authority

The introduction of AI / ML into critical workflows — such as estimating which molecules to pursue, which patient cohorts to recruit or which safety signals to prioritize — means IT is no longer just supporting the process; it becomes an integrated part of the decision chain.

- If a model cannot be exactly reproduced or if its training data cannot be traced and validated against ALCOA+ principles, the regulatory defensibility of the decision collapses.
- If an AI-assisted decision is later challenged by a regulator or a lawsuit, IT must defend the system and its outputs.

The moment AI influences a decision, IT becomes part of the decision chain.

This creates a new executive reality:

- If an AI-assisted decision is challenged, IT must defend the system
- If a model cannot be reproduced, the decision collapses
- If training data cannot be traced, the output is suspect

The most successful IT leaders do not sell cloud adoption. They design systems of control that enable innovation to move fast without breaking trust. AI, therefore, elevates IT from an enabling function to a decision-shaping authority, even if IT never touches the decision itself.

Growth through AI is only sustainable when:

- Models are governed
- Outputs are explainable
- Data is trusted
- Controls are auditable

This is why AI success in pharma is less about innovation speed and more about institutional trust.

IT is the steward of that trust.

Hybrid & Multi-Cloud are Not Preferences, But Risk Controls

For years, hybrid and multi-cloud strategies have been discussed as architectural styles, cost optimizations or vendor-negotiation tactics. In a regulated industry like pharma, hybrid and multi-cloud strategies are not merely architectural preferences driven by a desire for flexibility; they are non-negotiable strategies for risk containment, resilience and institutional trust.

At the executive level, the question is not "Should we be hybrid or multi-cloud?".

The real question is: "Where are we willing to accept single points of failure and can we defend that decision?"

Hybrid Architecture Protects What Cannot Fail or Be Moved

Pharma is not a greenfield industry. Certain systems have structural, not technical, constraints. Hybrid design is fundamentally a stability control that acknowledges the reality of pharma's structural constraints. Core Manufacturing Execution Systems (MES), long-standing validated GxP environments and systems tied to physical processes (such as specialized lab equipment) must remain close to the physical assets and stable for many years. Hybrid architecture allows the organization to preserve these validated, isolated environments, control the rate of change and selectively introduce cloud innovation (e.g., advanced analytics) without destabilizing critical operations. This approach represents risk-aware adoption.

Multi-Cloud is Resilience Strategy, not Feature Strategy

While secondary benefits include access to "best-of-breed services" and "avoiding vendor lock-in," the primary justification for multi-cloud is enterprise resilience and operational continuity. It significantly reduces dependency on a single provider's operational health, insulates the enterprise from region-wide failures (e.g., major outages in a single availability zone) and mitigates vulnerability to provider-specific policy, cost or technical changes.

Multi-cloud reduces:

- Dependency on a single provider's operational health
- Vulnerability to provider-specific policy or pricing changes
- Exposure to region-wide or service-wide failures
- Risk of systemic control-plane outages

Crucially, multi-cloud changes the shape and scale of failure. Instead of a catastrophic, enterprise-wide impact that halts manufacturing or clinical trials, failures become contained, localized, recoverable and most importantly, explainable to regulators. This ability to contain and articulate risk is the essence of advanced risk control.

Designing for Efficiency, Resilience, Control and Scalability

The pharmaceutical industry, driven by regulatory compliance, data-intensive research and the need for rapid innovation, faces unique challenges when designing its cloud strategy. The pharmaceutical industry, especially large US-based firms, is undergoing a digital transformation driven by cloud, data and AI. A modern cloud strategy for pharma must be cloud-agnostic and multi-cloud, leveraging best-of-breed services while avoiding vendor lock-in. The primary goals are to accelerate R&D, improve patient care and optimize commercial operations while ensuring compliance. A multi-cloud approach is increasingly adopted as the optimal path to meet stringent requirements for control, performance, scale, resiliency and cost efficiency.

A successful multi-cloud pharma strategy must be anchored in deliberate decisions across four key areas:

1. Cloud-Agnostic Data Strategy and Governance

Regulatory bodies such as the FDA require strict control over GxP (Good Manufacturing Practice, Good Clinical Practice, etc.) data. A cloud-agnostic data strategy ensures that critical data, while utilizing cloud services, remains under the organization's ultimate governance, independent of a single provider's infrastructure.

Strategy: Implement a robust Data Fabric or Data Mesh architecture that abstracts the underlying storage and processing layers. This approach allows data scientists and clinical teams to access data uniformly, regardless of whether it resides in AWS S3, Azure Blob Storage or Google Cloud Storage.

Key Enablers

- Unified Access Control: Centralized Identity and Access Management (IAM) across all clouds, often managed via an on-premises or cloud-native identity provider (e.g., Okta, Azure AD)
- Interoperable Data Formats: Enforcing open standards (e.g., Parquet, Avro) for storing analytical data to prevent vendor lock-in and facilitate seamless migration / replication
- Regulatory Zones: Designating specific cloud regions or accounts solely for GxP workloads, with enhanced audit logging and mandatory compliance controls (e.g., FedRAMP, HIPAA, GxP validation tools)

2. Hybrid and Multi-Cloud Networking and Connectivity

High-throughput, low-latency connectivity is crucial for R&D workloads, especially those involving genomic sequencing, molecular simulation and large-scale clinical data analysis. The networking strategy must ensure seamless, high-performance communication between on-premises systems, labs and diverse cloud environments.

Strategy: Utilize dedicated, secure interconnects (e.g., AWS Direct Connect, Azure ExpressRoute, Google Cloud Interconnect) coupled with a Software-Defined Wide Area Network (SD-WAN) overlay. This creates a unified "network backbone" that logically links all environments.

Key Enablers

- Global Load Balancing: Distributing application traffic across multiple cloud regions and providers to ensure high availability and disaster recovery
- Latency-Aware Routing: Employing intelligent routing protocols to direct high-priority computational jobs (e.g., drug discovery simulations) to the cloud provider offering the lowest network latency to the source data or HPC cluster
- Unified IP Management: Implementing a cohesive IP addressing scheme and next-generation firewall policies managed from a central point, simplifying security and compliance audits

3. Optimized Cloud Workload Placement

Not all workloads are created equal. The most cost efficient and performant strategy involves placing specific applications or data types in the cloud environment best suited for them, a principle often called “right-placing.”

Strategy: Categorize workloads based on their primary characteristics:

- SaaS / O365: Leverage vendor-specific environments for maximum collaboration (e.g., Microsoft Cloud for Healthcare)
- HPC or AI / ML: Place these in the cloud offering superior specialized hardware (e.g., specific GPU / TPU architectures) or specialized services (e.g., genomics sequencing services)
- Standard Enterprise Applications (ERP, LIMS): Utilize the cloud offering the most competitive utility pricing or where legacy infrastructure is already heavily invested (e.g., lift-and-shift to a primary cloud)
- Disaster Recovery (DR): Use a secondary, cost effective cloud provider purely for cold storage and non-production DR environments, leveraging competitive egress rates

Key Enablers:

- FinOps Culture: Implement real-time cost monitoring and optimization tools (e.g., CloudCheckr, Cloudability) across all providers to enforce spending guardrails and identify waste
- Containerization (Kubernetes): Use a portable orchestration platform (like Kubernetes / EKS / AKS / GKE) to move applications easily between clouds-based on performance needs or cost fluctuations

4. Unified DevSecOps and Automation Pipeline

Maintaining security and compliance across disparate environments is a major challenge. The solution lies in aggressive automation and a unified DevSecOps pipeline that enforces security controls before deployment.

Strategy: Build a “cloud-agnostic” automation layer using tools that manage infrastructure and configuration consistently across providers.

Key Enablers:

- Infrastructure as Codte (IaC): Use Terraform or Ansible to define infrastructure (VPCs, databases, security groups) once and deploy it across AWS, Azure and GCP, ensuring consistent configuration and compliance
- Centralized Security Policy Management: Employing tools like a Cloud Security Posture Management (CSPM) solution that continuously monitors security configuration and regulatory adherence (e.g., GxP audit trails) across all accounts, alerting or auto-remediating non-compliant resources
- Immutable Infrastructure: Deploy application environments using container images or machine images that cannot be modified after deployment, significantly reducing security drift

Data Strategy — The Intelligent Core

Pharma generates vast, heterogeneous data — from genomics and R&D to clinical trials and sales. An effective cloud data strategy unifies these silos into a cohesive, governed platform. A modern solution is a cloud data lakehouse that successfully combines the massive scale and low cost of traditional data lakes with the stringent governance, ACID properties (Atomicity, Consistency, Isolation, Durability) and schema enforcement of data warehouses, centralizing heterogeneous data (genomics, clinical, logs) into a cohesive, governed platform.

Cloud data platforms offer elasticity and global access: firms can store petabytes of genomic or clinical data and spin up on-demand compute for analysis. Pay-as-you-go pricing ensures you only pay for what you use — for example, running a large genomics pipeline for a few hours, then scaling down.

Key Tenets of Cloud Strategy for Enterprise Data Management and Analytics

Scalable and Modular Architecture
Cloud-native, modular architecture with scalability and adaptability to integrate data sources, AI services, models and frameworks easily.

Automation First
Automate pipelines, testing, deployment and monitoring (DataOps / MLOps / LLMOps) to improve reliability, speed and repeatability.

Data Security and Responsible AI by Design
Embed data security, privacy and ethical AI principles from the start.

Data Governance and Quality
Implement governance frameworks for data quality, lineage, privacy and compliance to ensure trusted and regulatory-ready data.

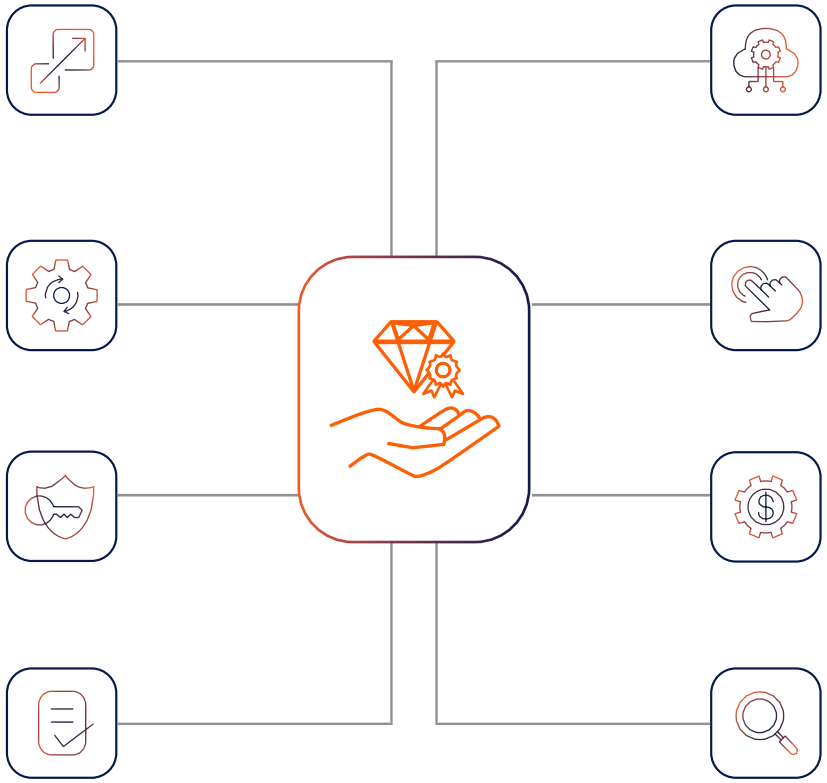

Data as a Product Mindset
Structure the platform around business domains, treat datasets and AI assets as products with clear business outcomes with measurable ROI.

Self-service and Democratization
Empower business users with self-service analytics, semantic layers and low / no-code tools while maintaining guardrails.

Cost Transparency and FinOps
Build in cost monitoring, chargeback models and workload optimization to control platform spend and prove value.

Observability and Reliability
Implement end-to-end observability for pipelines and models to monitor data drift, model performance and SLA adherence.

Building Blocks of Modern Analytics Architecture

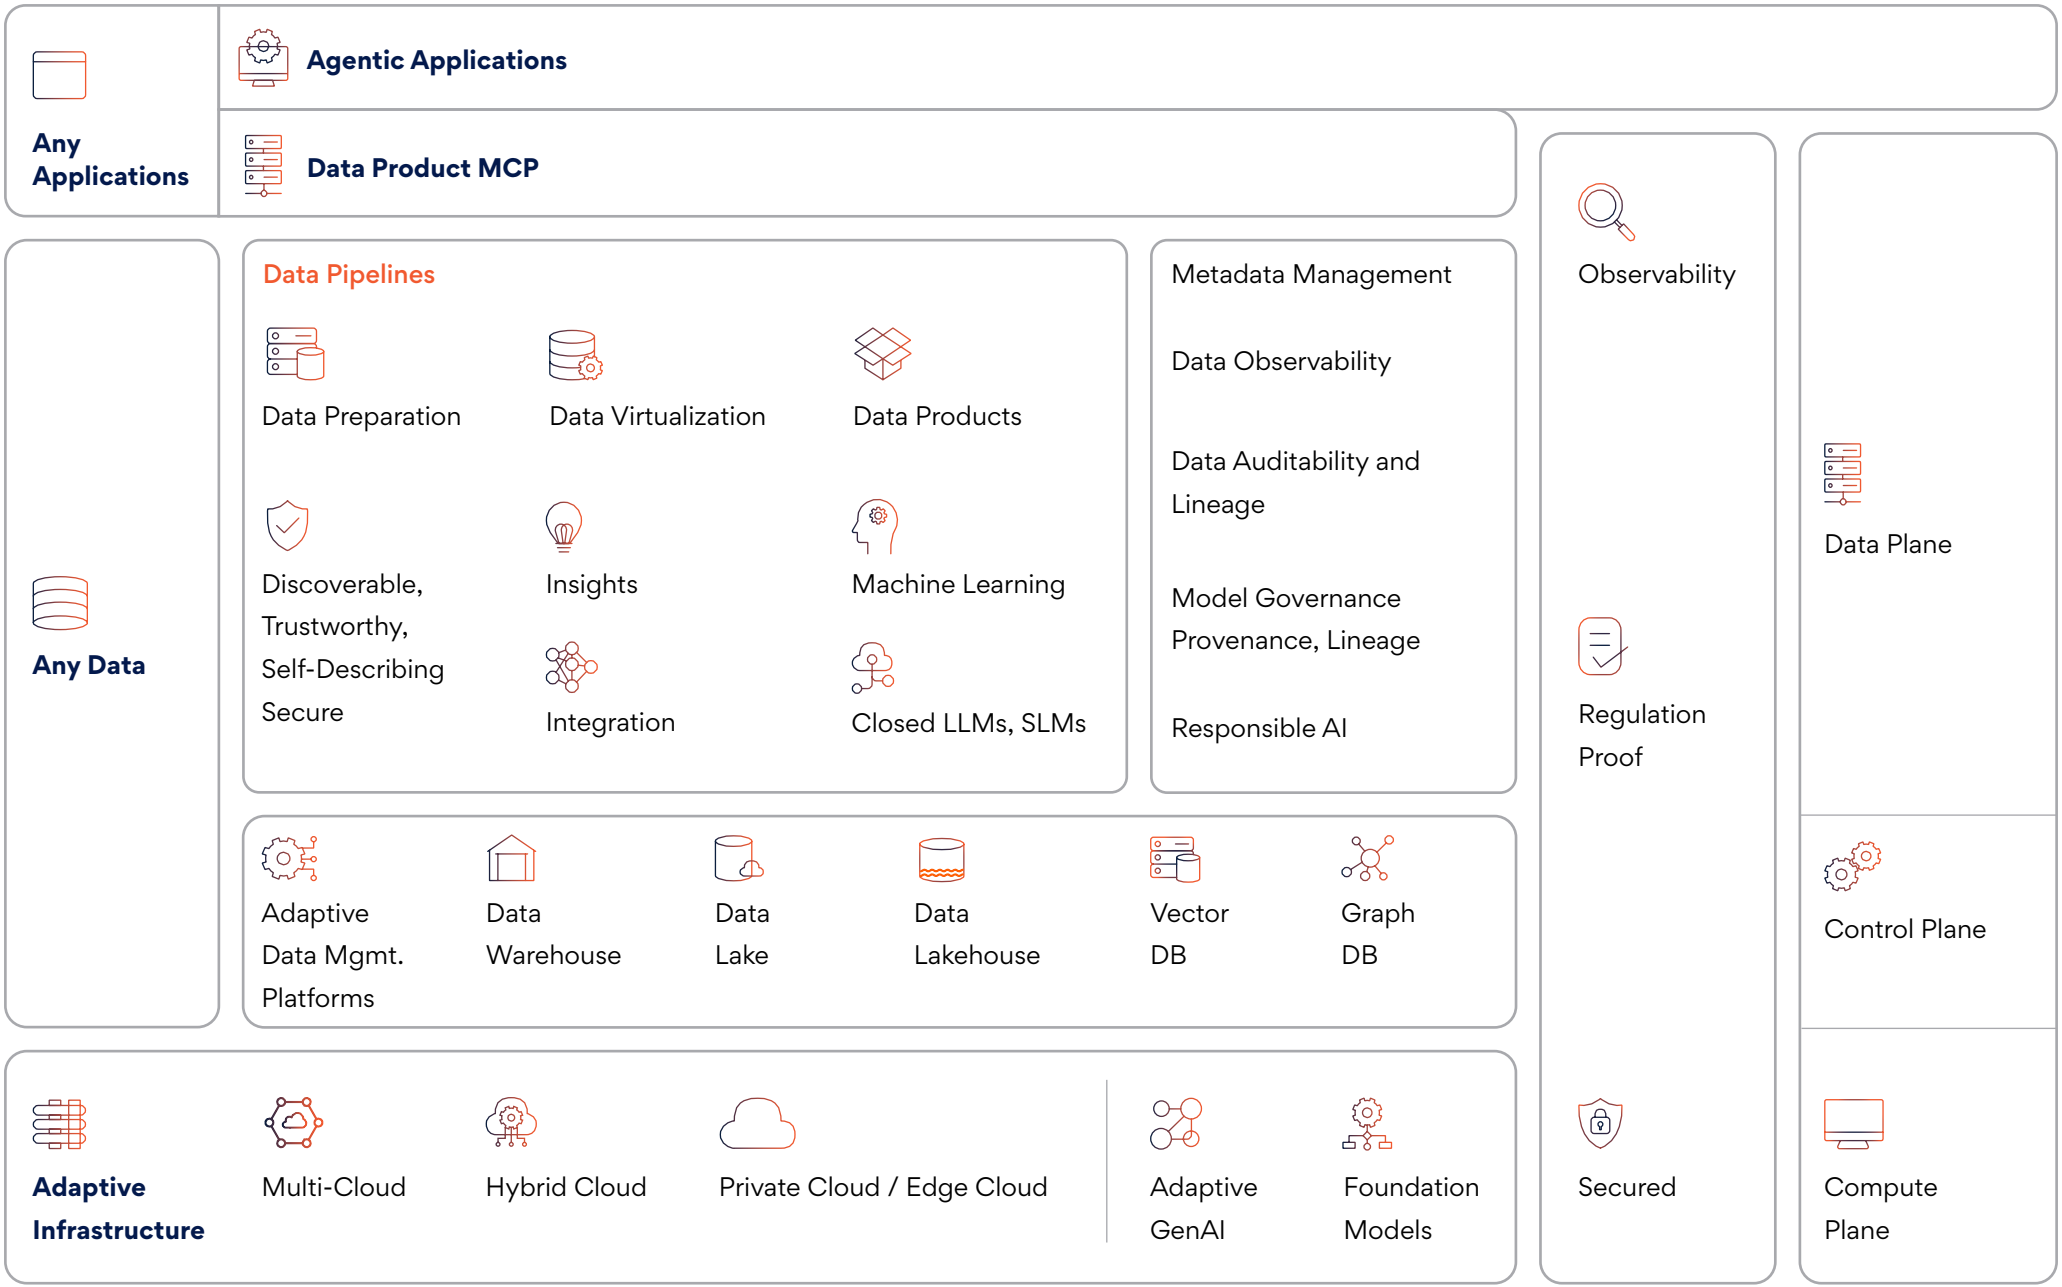

Data Platform Architecture for Enterprise Data Management and Analytics

Application Layer
BI / Data Visualization | Business Apps | Copilots | Agentic Applications

Access Layer
Data Products MCP | Integration

Any Data

Ingestion
Batch Ingestion
Streaming Ingestion
Near Real-time

Data Engineering and Processing
ETL / ELT Pipelines
Compute Engine

Master Data and Reference Data Management

AI / ML and GenAI
Feature Engg.
Model Training
Closed SLMs / LLMs
Fine-tuning
Vector Search

Agentic AI
Agent Frameworks / SDK
Design Patterns
Comm. Protocols

Data Management
Bronze → Silver → Gold

Semantic Layer
Business Abstractions
Harmonized KPI
Data Virtualization

Data Products
Market Place
Data Mesh
Catalog

Data Storage
Data Warehouse | Data Lake | Vector Database | Graph Database

Governance and Security Layer
Data Security at Rest and in-Motion
Role-based Access Controls
Data Auditability and Lineage
Business Glossary
Model Management
Technical Metadata
Business Metadata
Compliance

Data Observability
Data Quality
Data Reconciliation
Data Freshness

Platform Layer
CI / CD for Data, Analytics and AI
Infrastructure as Code
Performance Monitoring
Key Management
Workload Monitoring

Foundational Models
AI Gateway
Open LLMs

Architecture patterns should match needs. Many pharma organizations adopt a hybrid model: keep core, validated systems (like a certified clinical data warehouse) on-premises, while moving discovery and analytics workloads to the cloud. A common approach is to implement a data mesh or fabric for large enterprises, where domain teams (R&D, clinical, commercial) manage their own cloud data pipelines and tables but share a common governance layer. Over time, advanced firms layer in semantic models or knowledge graphs to enable easier access for scientists and analysts.

Data governance is critical. A formal governance framework defines who owns each data domain, access permissions, classification and quality rules. In pharma, this ensures every data item can be traced ("Who generated this clinical dataset? Was it modified?") — a necessity under FDA and HIPAA rules. Robust Master Data Management (MDM) ensures consistent identifiers (patients, trials, substances) across systems. Lineage and metadata catalog tools track where the data originated and how it been transformed. A good practice is to establish a governance committee and use automated pipelines with validation checks (e.g., flagging out-of-range values) built in.

The 2026 Architectural Shift: From Silos to Data Products

A cutting-edge data strategy moves past centralized warehousing toward a Data Mesh or Lakehouse architecture, where data ownership and quality are decentralized. Data is treated as a high-quality, consumable "product" owned by specific business domains (e.g., Genomics, Clinical Trials Operations).

- Federated Learning: Instead of pooling sensitive patient data, pharma companies are increasingly using federated learning. This allows them to train AI models across different regions or even different companies without ever moving the raw data out of its original secure environment (e.g., keeping EU data in the EU to satisfy GDPR).
- Digital Twins: For R&D, companies are deploying Digital Twins of clinical trials. By simulating patient behavior and drug-target interactions in a cloud-based "in silico" environment, researchers can predict protocol failures before a single patient is enrolled.

The Core Infrastructure: AI-Native Data Lakehouse

The modern standard is a cloud-native lakehouse (e.g., Databricks on Azure / AWS or BigQuery on Google Cloud) that merges the massive scale of a data lake with the governance of a warehouse.

- Federated Data Fabric: Instead of moving petabytes of genomic data, 2026 strategies use a data fabric to query data in place (e.g., using Starburst / Trino). This is vital for international trials where data residency laws (GDPR, etc.) prevent moving patient data across borders
- Vectorized Multi-Modal Search: Research now involves text (medical notes), images (pathology slides) and sequences (genomics). Strategic workloads now use vector databases (such as Pinecone or PgVector) to allow researchers to search for "patients with similar T-cell exhaustion markers" across disparate datasets

GxP Compliance-as-Code (CaC)

Modern strategies automate the validation of cloud infrastructure. CaC has evolved from a niche DevOps trend into the mandatory operating model for pharma cloud deployments. It replaces the traditional paper-and-ink validation process, which took months, with automated, version-controlled scripts that prove a system is in a validated state in real-time. It is recommended to use CI / CD pipelines (GitHub Actions / Azure DevOps) to automatically generate the documentation required by the FDA / EMA whenever a research workload is updated. This ensures the research environment is reproducible.

- Validated Modules: Create gold images of cloud environments that are pre-approved by the Quality Assurance (QA) team
- Drift Detection: If a researcher accidentally changes a security setting (e.g., making a bucket public), the code automatically detects this drift from the validated state and rolls it back instantly
- Traceability Matrix: The system automatically generates a digital Traceability Matrix, linking User Requirements (URS) to Functional Specs (FS) and, finally, to the automated test results
- Digital Signatures: Approvals are captured via automated e-signature workflows (21 CFR Part 11 compliant) within the pipeline itself

AI Transformation in Pharma

AI, especially Generative AI (GenAI), is poised to revolutionize pharma across R&D, clinical development, manufacturing and commercial functions. Studies show nearly all aspects of the industry will be affected, from faster drug discovery to ultra-targeted marketing.

For example:

- Drug Discovery: Generative models and GPT-like tools can "extract scientific knowledge" from vast literature and patents to highlight new targets. By alleviating the manual literature review, such tools can increase the initial identification of drug targets by over 30%. Similarly, AI-driven chemistry models can screen millions of compounds in silico, predicting promising candidates weeks faster than traditional labs. McKinsey reports that AI-enabled screening can boost model performance by ~2.5× and cut lead identification time by >4×. In short, accelerated R&D pipelines reduce time and cost to find novel therapies.
- Clinical Development: AI can optimize trial design and operations. GenAI assistants (or co-pilots) help with patient selection by analyzing genetic and real-world data, ensuring more representative cohorts. They can auto-draft regulatory submissions and study protocols. For instance, AI can suggest relevant references during document writing, thereby speeding reviews. Industry analysis estimates these tools could reduce clinical trial duration by up to 1–2 years and cut costs by ~20%, effectively doubling a project's NPV. One use case found patient-matching AI raising trial success probability by 10% while slashing cost / duration by 20%.
- Manufacturing and Supply Chain: AI-driven forecasting and optimization yield big gains. GenAI agents can propose sourcing options, optimize batch processes and foresee inventory shortages. For example, no-touch planning assistants continuously rebalance raw material orders in real-time. Enhanced predictive maintenance on manufacturing equipment is another major use; machine learning models trained on sensor data can pre-emptively flag quality issues. While less mature, such applications will improve yield and responsiveness.
- Commercial and Marketing: AI will reshape engagement. Generative models can create personalized content rapidly: pharma marketers have piloted AI-powered creative that drafts compliant promotional materials in days (versus weeks). Early results show 30–50% cuts in content creation time and cost. AI also augments field teams: sales co-pilots can synthesize customer insights on-demand, helping reps tailor their pitches. Studies suggest this could boost sales productivity 10–15%, translating into ~1–2% topline growth. Even medical affairs stand to benefit, with AI summarizing key opinion leader insights and automating literature review (though with care around hallucination risk).

AI's potential value is enormous (McKinsey estimates a potential impact of $60–110 billion annually by 2030), but successful deployment is entirely contingent on data readiness, quality and governance. The immediate focus must be on securing and validating data pipelines, investing heavily in Machine Learning Operations (MLOps) for model lifecycle management and establishing transparent AI ethics and governance frameworks. The end vision is integrated cloud data platforms and AI co-pilots that accelerate every critical stage of discovery, compliance and commercial success.

Conclusion

A successful cloud and AI strategy for the modern pharmaceutical company is not a project list; it is a comprehensive, multi-layered roadmap that fuses technology with regulation.

- **Infrastructure and Resilience**: Architect a robust, multi-cloud, high-availability foundation with rigorous security and automated GxP / data privacy compliance controls enforced at the architecture level
- **Finance and Optimization**: Implement a disciplined FinOps culture and practice to transparently track and optimize costs across all cloud providers and business units, ensuring sustainable growth
- **Data as Strategy**: Build a modern, governed data platform (lakehouse / mesh) to unify R&D, clinical and commercial data, transforming data from a liability into a highly trusted, reusable asset
- **Innovation and Trust**: Leverage AI across the entire value chain, from molecular discovery to market operations, with an uncompromising, auditable focus on data quality, model trust and seamless user adoption

To know how Persistent can help pharma organizations choose the right cloud strategy, reach us at **grow_gcp@persistent.com**

About Persistent

Persistent Systems (BSE: 533179 and NSE: PERSISTENT) is a global services and solutions company delivering AI-led, platform driven Digital Engineering and Enterprise Modernization to businesses across industries. With over 26,500 employees located in 18 countries, the Company is committed to innovation and client success. Persistent offers a comprehensive suite of services, including software engineering, product development, data and analytics, CX transformation, cloud computing and intelligent automation. The Company is part of the MSCI India Index and is included in key indices of the National Stock Exchange of India, including the Nifty Midcap 50, Nifty IT and Nifty MidCap Liquid 15, as well as several on the BSE such as the S&P BSE 100 and S&P BSE SENSEX Next 50. Persistent is also a constituent of the Dow Jones Sustainability World Index. The Company has achieved carbon neutrality, reinforcing its commitment to sustainability and responsible business practices. Persistent has also been named one of America's Greatest Workplaces for Inclusion & Diversity 2025 by Newsweek and Plant A Insights Group. As a participant of the United Nations Global Compact, the Company is committed to aligning strategies and operations with universal principles on human rights, labor, environment and anti-corruption, as well as take actions that advance societal goals. With 468% growth in brand value since 2020, Persistent is the fastest-growing IT services brand in 'Brand Finance India 100' 2025 Report.

USA
Persistent Systems, Inc., 2055 Laurelwood Road, Suite 210, Santa Clara,
CA 95054, Tel: +1 (408) 216 7010, Fax: +1 (408) 451 9177,
Email: info@persistent.com

India
Persistent Systems Limited, Bhageerath, 402, Senapati Bapat Road,
Pune 411016, Tel: +91 (20) 6703 0000,
Fax: +91 (20) 6703 0008

Persistent
Re(AI)magining™ the World