IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Roche NVIDIA AI Factory: Digital Twins for GLP-1 Pharma

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-04-14 · roche nvidia ai · glp-1 manufacturing · pharma digital twins · ai drug discovery · nvidia omniverse · biopharma infrastructure · lab-in-the-loop

Original article: https://intuitionlabs.ai/articles/roche-nvidia-ai-factory-glp-1-digital-twins



In this article

1. 1. Introduction
2. 2. Roche’s NVIDIA AI Factory Expansion
3. 3. Digital Twins & Omniverse in Pharma Manufacturing
4. 4. Case Studies and Industry Context
5. 5. Data Analysis & Evidence
6. 6. Discussion and Future Directions
7. 7. Conclusion

Executive Summary: Roche – the global biotech and diagnostics leader – has dramatically scaled its AI and computing infrastructure through a new partnership with NVIDIA. In March 2026, Roche announced the addition of **2,176 NVIDIA Blackwell GPUs** on premises across the United States and Europe, bringing its combined on-premises and cloud infrastructure to over **3,500 Blackwell GPUs**; Roche described this as the largest announced GPU footprint available to a pharmaceutical company (Roche announcement) (www.biospace.com) (www.datacenterdynamics.com). This "AI factory" is powering Roche's entire value chain – from discovery ("Lab-in-the-Loop" model for molecular design) to diagnostics and manufacturing – by embedding AI deep into drug R&D workflows and plant operations (blogs.nvidia.com) (www.biospace.com). Key NVIDIA technologies include the BioNeMo foundation-model platform (to accelerate molecular design), Parabricks (for secondary analysis of DNA and RNA sequencing data), NeMo/Guardrails (for safe generative AI), and especially Omniverse (NVIDIA's simulation platform) to create **high-fidelity digital twins** of production lines (blogs.nvidia.com). In particular, Roche is using Omniverse to **simulate its new GLP-1 (incretin) manufacturing plant in Holly Springs, NC**, to help accelerate facility development (blogs.nvidia.com) (www.biospace.com).

This initiative is timely: the global market for GLP-1 obesity and diabetes drugs is exploding (projected ~USD 62 billion in 2025 to ~USD 157 billion by 2035 (www.globenewswire.com), with forecasts of ~$200 billion by 2030 (www.jpmorgan.com) and tens of millions of patients on GLP-1s by the end of the decade). Roche itself has rapidly invested in this space – for example acquiring Carmot Therapeutics (gaining a pipeline of GLP-1/GIP agonists and an oral GLP-1) and building a new $2 billion NC facility that will "support production of next-generation treatments for obesity" (www.roche.com) (www.biospace.com). McKinsey describes potential gains in quality, speed, agility and resilience from digital and analytics advances in biopharma manufacturing (McKinsey analysis). By uniting vast AI compute ("AI Factory") with robotic labs and digital simulation, Roche aims to **compress drug development cycles**, optimize manufacturing throughput, and accelerate delivery of life-changing medicines.

This report provides an in-depth analysis of the Roche-NVIDIA AI factory and its implications for GLP-1 drug manufacturing. We begin with background on Roche, GLP-1 therapeutics, and digital twin technology, and then detail Roche's AI infrastructure expansion and NVIDIA's technologies. We examine how Omniverse-driven digital twins are applied to Roche's NC plant design, review case studies (e.g. competitors like Eli Lilly's LillyPod supercomputer (www.techtarget.com) and other Omniverse factory uses), and present data on AI-driven efficiency gains. Finally, we discuss challenges, broader industry trends, and future directions in AI-powered pharma manufacturing. All claims are supported by extensive citations.

1. Introduction

1.1 Roche in the BioPharma Industry

Roche (F. Hoffmann-La Roche Ltd) is a Swiss company founded in 1896, now one of the world's largest biotechnology and diagnostics firms (www.biospace.com). It operates two main divisions: Pharmaceuticals (drugs and biologics) and Diagnostics (lab testing devices). Roche has pioneered targeted therapies and personalized medicine, especially in oncology and immunology, and has historically invested heavily in R&D. For 2025, Roche reported **CHF 12.2 billion in core R&D investments** and **CHF 61.5 billion in Group sales** (Roche Annual Report 2025), equivalent to approximately **19.8%** on that basis.

In recent years, Roche has also embraced digitalization across its businesses. Its Diagnostics arm processes millions of patient samples daily, requiring sophisticated data pipelines. Roche's Pharma arm has pursued innovations like *Lab-in-the-Loop* – an integrated approach linking automated biology/chemistry experiments with AI-driven hypothesis testing to accelerate discovery (blogs.nvidia.com) (www.biospace.com). Roche's leadership emphasizes that "time is the most critical variable" in healthcare, noting that every day saved in development means a treatment reaches patients sooner (www.biospace.com). Thus, the company views AI and computing power as core tools to reduce lead times across R&D, clinical trials, and manufacturing.

1.2 GLP-1 Therapeutics and Market Dynamics

Glucagon-like peptide-1 (GLP-1) receptor agonists are a class of peptide hormones originally developed to treat type 2 diabetes. GLP-1 drugs (e.g. liraglutide, semaglutide, tirzepatide) mimic a gut hormone that stimulates insulin secretion and suppresses appetite. In the past few years, several GLP-1 drugs have shown dramatic weight-loss effects, not only controlling blood glucose but also shrinking adipose tissue. This has generated enormous demand: patients and prescribers use GLP-1s for obesity management and related metabolic disorders.

The GLP-1 market has become a **multibillion-dollar global opportunity**. Research forecasts project the GLP-1 market to grow from around **USD 62 billion** in 2025 to **~USD 157 billion by 2035** (CAGR ~9.7%) (www.globenewswire.com), with North America and Asia leading growth.JPMorgan recently estimated that approximately **25 million Americans could be on GLP-1 treatment by 2030**, up from about 10 million in 2025 (and only ~5 million in 2023) (www.jpmorgan.com). These figures reflect both increasing treatment of diabetes patients and the huge surge in prescribing GLP-1s for obesity. In the United States, FDA labeling for oral Wegovy (semaglutide) includes long-term weight reduction in adults with obesity, or overweight with at least one weight-related comorbid condition, in combination with a reduced-calorie diet and increased physical activity (FDA prescribing information). The sheer scale of demand has strained supply chains: even market leaders like Novo Nordisk (Wegovy, Ozempic) and Eli Lilly (Mounjaro, Zepbound) have increased production capacity aggressively. For example, Novo announced it would open a new major plant for Wegovy by 2027 and lower pricing to boost volume (www.novonordisk.com).

Roche, historically not a GLP-1 player, has aggressively entered this arena. In December 2023 Roche agreed to acquire Carmot Therapeutics for **USD 2.7 billion upfront** to gain Carmot's portfolio of GLP-1/GIP agonists and an oral GLP-1 candidate (www.fiercebiotech.com). The Carmot-derived candidates have different development paths. Roche's June 2026 update described enicepatide (CT-388) moving into Phase III development (Roche update). In its July 2026 presentation, Roche stated that development of acmopatide (CT-868) would be discontinued (Roche half-year presentation, slide 48). CT-996 is being investigated as an oral, once-daily treatment for obesity and type 2 diabetes (Genentech). In January 2024 the Carmot deal closed, bringing Roche into position as a contender behind Novo and Lilly.

To support this program at scale, Roche is building manufacturing capacity in the United States. On 25 August 2025 Roche's Genentech unit broke ground on a **$700 million** biomanufacturing facility in Holly Springs, North Carolina (www.fiercepharma.com) (www.roche.com). This plant – Genentech's first East Coast site – is explicitly designed for "next-generation metabolic medicines, including treatments for obesity" (www.roche.com). It will be a high-volume, automated site (built with advanced digital controls and sustainability in mind) and is planned for operation by 2029 (www.roche.com) (www.biospace.com). In early

2026 Roche more than **doubled** its investment to USD 2.0 billion (adding another $1.3 billion) for this facility, reflecting the urgency of the obesity-drug race (www.biospace.com). At full build-out the facility will support Roche's GLP-1 and related metabolic drug systems, performing **fill/finish for injectable therapies** (www.fiercepharma.com) and potentially even producing active ingredients.

1.3 Artificial Intelligence in Pharma R&D and Manufacturing

The biopharma industry has seen a "digital transformation" over the last decade. High-throughput screening, computational chemistry, and genomics have long relied on supercomputers. However, recent advances in machine learning – especially deep learning – are accelerating this trend. Pharmaceutical companies increasingly apply AI to: predict protein structures (e.g. AlphaFold), design novel molecules, optimize formulation, and even plan clinical trials. For instance, Roche's Genentech has pioneered "Lab-in-the-Loop" processes, where human scientists, lab automation, and AI models form a closed feedback cycle (blogs.nvidia.com) (www.biospace.com). According to Roche, nearly **90% of Genentech's eligible small-molecule programs already integrate AI** into design or analysis (blogs.nvidia.com).

Industry-wide, big pharma and biotech have in recent years announced major AI or supercomputing initiatives. For example, in January 2026 NVIDIA and Eli Lilly announced an AI co-innovation lab, committing jointly to invest up to $1 billion over five years in talent, infrastructure and compute (NVIDIA announcement). Other players like AstraZeneca and Sanofi likewise report extensive machine learning in their pipelines. The motivation is clear: drug discovery is lengthy and costly (often >10 years to bring a new drug to market), so reducing any step – target identification, lead optimization, toxicity screening – can save years and costs.

Similarly, manufacturing in pharma stands on the cusp of an Industry 4.0 revolution. Digital twins – high-fidelity virtual models of physical plants or processes – are becoming feasible with modern simulation and sensor networks. In other industries (automotive, electronics, semiconductors), digital twins have shown their power: for example, Foxconn uses NVIDIA Omniverse-based twins to simulate entire factory layouts and robotic cell operations (blogs.nvidia.com), achieving faster commissioning of new lines. In pharma, digital twins promise to optimize bioreactors, fluid flows, filling lines, and complex automation sequences. The idea is to **simulate a production line end-to-end** (including chemical processes, robotics, equipment faults, etc.) so that engineers can optimize performance and avoid costly rework. McKinsey describes potential gains in quality, speed, agility and resilience from digital and analytics advances in biopharma manufacturing (McKinsey analysis).

For example, Cambridge University and Singapore's A*STAR recently developed an AI-driven digital-twin platform for pharma plants. Their system uses real-time sensor streams and ontologies to create "virtual replicas" of equipment, enabling anomaly detection (e.g. mismatched flow rates) and predictive maintenance through machine learning (www.europeanpharmaceuticalreview.com). The researchers note that such twins can embed expert knowledge and potentially support scheduling and quality monitoring (www.europeanpharmaceuticalreview.com). Industry publications echo this: *PharmaFocus Europe* declares that "Digital Twins in Pharma are revolutionizing drug development, bioprocessing, and personalized medicine" by enabling dynamic optimization and regulatory-ready modeling (www.pharmafocuseurope.com).

There are still hurdles – validating AI models against regulatory standards, integrating legacy equipment, and ensuring data quality – but the momentum is strong. In practical terms, a pharma company that can digitally iterate its lab processes and factories (instead of physical trial-and-error) may achieve breakthroughs that

translate to faster time-to-market and lower costs. Disciplines like Process Analytical Technology (PAT) and advanced process control (APC) are becoming integrated into these twins to ensure compliance. For instance, digital twin environments allow manufacturers to store and digitize validation data, creating an audit trail for quality systems (www.europeanpharmaceuticalreview.com). As one industry analyst put it, companies that "scale AI will define the next wave of human progress" – a view echoed by Siemens and others advocating for scaling AI into production lines (the so-called "AI factory") (pradeepstechpoints.wordpress.com).

1.4 NVIDIA's Role in AI and Digital Manufacturing

NVIDIA is widely recognized as a leader in AI hardware and software. Its GPUs (Graphics Processing Units) are the de facto standard accelerators for deep learning and scientific computing. The **NVIDIA Blackwell** architecture – the successor to Hopper – underpins the latest high-end GPUs in production today. A Blackwell Ultra GPU, for example, contains **160 streaming multiprocessors (SMs)** across a dual-die design and packs **288 GB of HBM3E memory** (developer.nvidia.com). It introduces a new FP4 precision (NVFP4) that achieves up to **15 petaFLOPS** of dense AI compute per chip, along with full NVLink connectivity and very high bandwidth (developer.nvidia.com). In practical terms, a DGX SuperPOD made of dozens of such GPUs can deliver exascale-class AI performance (hundreds of PFLOPS) for training massive models.

But NVIDIA's contribution goes beyond raw silicon. The company also provides **full-stack software and platform capabilities** for live data analytics, digital twin simulation, and model generation. Key offerings include:

- **NVIDIA BioNeMo:** An end-to-end platform and library for "biological neural modeling" (foundation models in biology). It enables the training of large language or transformer models on proteomics, genomics, and chemistry data (blogs.nvidia.com). BioNeMo is tailored for drug discovery tasks such as protein folding inference or molecular generation, and it ties into Roche's Lab-in-the-Loop approach (blogs.nvidia.com).
- **NVIDIA Isaac Sim and Simulation:** Formerly known as Isaac, this suite provides robotics simulation. Combined with Omniverse, it lets engineers simulate robotic work cells and automation in a virtual version of the factory. (E.g. Foxconn and Wistron use Isaac + Omniverse to train assembly robots before deployment (blogs.nvidia.com).)
- **NVIDIA Parabricks:** GPU-accelerated software for secondary analysis of next-generation DNA and RNA sequencing data, rather than pathology-image segmentation (NVIDIA documentation).
- **NVIDIA NeMo and NeMo Guardrails:** NeMo is a framework for building and deploying large language models; "Guardrails" are extensions that enforce safety and reliability in conversational AI. Roche uses NeMo Guardrails in its digital health tools (blogs.nvidia.com).
- **NVIDIA Omniverse:** A real-time 3D simulation and collaboration platform. It uses **Universal Scene Description (USD)** to link CAD, physics, and AI agents in one virtual environment. Omniverse allows multiple engineers (even around the world) to work on the same simulated model simultaneously. It provides libraries for physics simulation, robotics (via Isaac), visual rendering, and easily integrates with tools like Siemens NX, CAD packages, and factory process models.

NVIDIA often refers to its concept of an **“AI Factory”** – the idea of an end-to-end AI-accelerated workflow for an enterprise. While the term is used broadly (even for data centers run by AWS/Microsoft (pradeepstechpoints.wordpress.com)), in Roche’s context it means a centralized HPC/cloud environment that *powers* discovery, diagnostics, and manufacturing through AI models and simulations. To this end, Roche is setting up a **hybrid cloud/on-prem lens**: high-performance supercomputing units (DGX or similar) at Roche’s campuses, plus access to NVIDIA’s DGX Cloud or other hyper-scale cloud GPU pools. This architecture allows massive grid-scaler training (for example, training a multi-billion-parameter molecule predictor) while also enabling site-local development (teams working on a protected pipeline with their own GPU cluster) (blogs.nvidia.com) (www.datacenterdynamics.com).

Table 1. NVIDIA Platforms and Roche Use Cases. Below is an overview of key NVIDIA technologies in Roche’s AI factory and how Roche plans to use them:

NVIDIA Technology	Function	Roche Use Case
NVIDIA Blackwell GPUs	Accelerated computing	Roche announced 2,176 added on-premises GPUs across the United States and Europe and a combined on-premises/cloud footprint exceeding 3,500 Blackwell GPUs (Roche announcement).
NVIDIA BioNeMo	Foundation-model platform for biology	Enhances Roche’s Lab-in-the-Loop drug discovery: train molecular and biological AI models, generate novel compound ideas, and analyze experimental data at scale (blogs.nvidia.com).
NVIDIA Omniverse	3D simulation/collaboration platform	Build high-fidelity *digital twins* of manufacturing plants (e.g. Holly Springs GLP-1 facility) to simulate layouts, robotic cells, and processes to help accelerate facility development (blogs.nvidia.com) (www.biospace.com).
NVIDIA Parabricks	GPU-accelerated genomics	Secondary analysis of next-generation DNA and RNA sequencing data (NVIDIA documentation).
NVIDIA NeMo & Guardrails	Generative AI and configurable controls	Roche uses NeMo Guardrails in digital health. Fact-checking guardrails can help reduce hallucinations and false claims; this is a capability, not a guarantee of deployment-wide compliance, accuracy or patient safety (NVIDIA documentation).
Metropolis & Video Insight	Video AI library	Potential future use for automated laboratory monitoring or facility security (Roche has not specifically reported use yet).

2. Roche’s NVIDIA AI Factory Expansion

2.1 Deployment of GPUs and Architecture

In March 2026, Roche formally announced its global AI infrastructure expansion in partnership with NVIDIA (www.biospace.com). This builds on a 2023 collaboration (Genentech–NVIDIA) where they optimized machine learning algorithms on NVIDIA hardware (www.datacenterdynamics.com). The centerpiece is **2,176 new NVIDIA Blackwell GPUs deployed on-premises** across sites in the United States and Europe (www.biospace.com). Combined with existing capacity and cloud resources, Roche now has **over 3,500**

NVIDIA GPUs available enterprise-wide (all the GPUs are Blackwell generation) (www.biospace.com) (www.datacenterdynamics.com). Roche claims this is the largest announced GPU footprint in the pharmaceutical industry to date (www.biospace.com) (www.datacenterdynamics.com). (By comparison, rival Eli Lilly's recently unveiled supercomputer "LillyPod" has 1,016 GPUs (www.techtarget.com).)

The hybrid architecture – mixing on-premise clusters and cloud – is intentional. On-premises clusters at Roche sites (likely linked via NVSwitch and NVLink fabrics internally) handle sensitive proprietary workloads and support hybrid data flows, while public cloud GPU pools provide virtually unlimited scaling when needed. Roche notes this lets teams "train massive models" in the supercomputing portion **while supporting local development across global sites** (blogs.nvidia.com). In effect, bench scientists and data scientists at different locations can leverage the same underlying AI platform. The company has emphasized that this AI factory is an *enterprise-grade* backbone, embedding AI "from discovery to development, manufacturing and commercialization" (www.biospace.com). Wafaa Mamilli, Roche's Chief Digital & Technology Officer, described it as combining "world-class computing power with Roche's scientific expertise" to bring AI to every part of the drug and diagnostic pipeline (www.biospace.com) (blogs.nvidia.com).

Table 2. Roche AI Factory Scale (Approximate)

Aspect	Figures / Notes	Source
Added on-premises GPUs	2,176 NVIDIA Blackwell GPUs across the United States and Europe.	Roche announcement
Combined GPU footprint	More than 3,500 Blackwell GPUs across on-premises and cloud infrastructure.	Roche announcement
Software and applications	BioNeMo for Lab-in-the-Loop research; Omniverse for manufacturing digital twins; Parabricks for data analysis; NeMo Guardrails in digital health.	Roche announcement

2.2 AI-Powered R&D Workflow ("Lab-in-the-Loop")

Even before this expansion, Roche had long pursued AI in early R&D. Roche and its Genentech subsidiary pioneered the idea of a *"Lab-in-the-Loop"*, where experiments feed data into AI models which then suggest the next experiments in an iterative cycle (blogs.nvidia.com). With the new GPU supercomputer, this Lab-in-the-Loop strategy can be scaled across many projects concurrently. Roche reports that in drug discovery applications "nearly 90% of Genentech's eligible small-molecule programs integrate AI" (blogs.nvidia.com). For example, generative AI has already enabled specific projects: one oncology **degrader molecule** was designed 25% faster with AI assistance, and in another program a backup drug candidate was discovered in 7 months instead of 2+ years (blogs.nvidia.com). These impressive anecdotes highlight how more compute can collapse development timelines for medicinal chemistry.

On the platform side, Roche is using **NVIDIA BioNeMo**, which is optimized for biological sequence data and molecule generation. With BioNeMo running on Roche's Blackwell GPUs, scientists can train or fine-tune large language models (LLMs) or diffusion models on protein/genomics/compound data. For instance, one could

train a transformer model on millions of protein sequences to predict structural features or binding affinities, or train a generative model to propose novel drug-like molecules. Roche specifically notes that BioNeMo "valuable for molecular and foundation models", enabling exploration "far larger swaths of the biological and chemical space" at high speed (blogs.nvidia.com). In practice, Roche is likely integrating BioNeMo with its internal chemical databases, assay results, and high-throughput screening data to automate hypothesis generation.

2.3 Applications in Diagnostics and Clinical Data

While primarily a pharma initiative, the AI factory also benefits Roche's diagnostics business. Roche's diagnostics division operates large clinical labs and diagnostic devices (e.g. blood analyzers, imaging scanners, sequencing equipment). These generate diverse data (images, signals, omics). With GPU acceleration, Roche can run advanced analysis at scale.

In digital pathology, for example, Roche uses NVIDIA acceleration to **scan thousands of pathology images**, looking for disease patterns that might elude the human eye (blogs.nvidia.com). Parabricks – originally developed for genomics – has been repurposed or analogously used to rapidly align and analyze sequencing data from diagnostic tests, as hinted by Roche's emphasis on "insights across vast datasets" (blogs.nvidia.com). In digital health, Roche uses **NVIDIA NeMo Guardrails** (blogs.nvidia.com).

The combined effect is that both therapeutic R&D and diagnostics inform one another. For example, large language models trained on electronic health records or scientific literature (hosted on Roche's GPUs) could identify new target indications or safety signals. The press materials hint that Roche is "scanning a large number of images" and using AI to provide clinical decision support, though details remain proprietary (blogs.nvidia.com). What is clear is that Roche envisions this AI factory as **enterprise-wide**: "shifting AI from a specialized discipline to a core capability embedded in everyday workflows," from lab research to clinical diagnostics to manufacturing quality control (blogs.nvidia.com).

2.4 Roche's AI-Ware Empowering Manufacturing and Quality

Roche explicitly regards manufacturing and supply-chain as critical targets for AI. Small efficiency gains in a plant can scale into major global benefit (improved yield, fewer batch failures, faster new-product ramp-up). The press release states that even functions like regulatory documentation, quality assurance, and production scheduling are being enhanced with AI (blogs.nvidia.com). For example:

- **Quality Control:** AI vision systems (perhaps NVIDIA Metropolis) could inspect vials and reagents for microdefects, leveraging Omniverse-simulated images for training datasets. In one section on diagnostics, Roche noted that 3D digital pathology scanning is more efficient with GPU acceleration (blogs.nvidia.com); a similar idea applies in QC for manufacturing.
- **Scheduling and Supply:** Complex biologics manufacturing can be thought of as a sequence of bioreactors, centrifuges, filtrations, and filling lines. AI-driven scheduling – which takes into account equipment availability, turnover times, and regulatory testing – can optimize batch queuing. Though not detailed in the announcements, such work is within scope for an "AI-accelerated healthcare organization".

Importantly, the NVIDIA Omniverse platform underlies much of this. Omniverse allows creation of detailed 3D models of entire production facilities (walls, machines, conveyors, AGVs, robots, even staff movement). These digital twins can simulate every aspect of a plant. For Roche's GLP-1 facility in North Carolina, NVIDIA reports that Omniverse-based digital twins are helping accelerate facility development by simulating and optimizing systems before they go live. The announcement does not establish deployment of the specific clean-in-place and detailed material-flow workflows discussed here (NVIDIA announcement). As Wafaa Mamilli put it: "We are using NVIDIA Omniverse to build high-fidelity digital twins of our manufacturing sites such as the [Holly Springs] facility in North Carolina" (blogs.nvidia.com). This approach mirrors what electronics manufacturers are doing (e.g. Foxconn's Fii Digital Twin platform, built on Omniverse, has optimized entire production lines (blogs.nvidia.com)) and what semicon fabs are doing (e.g. TSMC converting 2D designs into 3D twin simulations to pre-empt layout collisions (blogs.nvidia.com)).

Because Roche's NC site is still under design/construction, the Omniverse twin can continuously incorporate design changes, updated equipment placements, and new automation scripts. This virtual plant will link to real-time engineering data. When the physical plant is built, the digital twin can then be used in parallel to monitor operations (a concept known as a "living twin"): sensor streams from the real plant (temperatures, pressures, flow rates) could be fed into the twin, enabling early detection of drift or faults. In effect, this closes the loop: the same simulation that helped design the plant will eventually help run it.

Illustration of how Omniverse-based digital twins improve plant design: Imagine modifying the piping layout in the digital 3D model and instantly checking for flow imbalances or maintenance access, *without* having to physically reconfigure anything. Nvidia sees such use cases as central to "Industrial Metaverse" initiatives – factory digital twins that are integrated with AI and robotics to maximize efficiency (blogs.nvidia.com). Roche's adoption of Omniverse for its GLP-1 line places it at the forefront of this trend in pharma manufacturing.

3. Digital Twins & Omniverse in Pharma Manufacturing

3.1 What Are Digital Twins and Why They Matter

A *digital twin* is a dynamic virtual replica of a physical system, ranging from individual devices to entire factories (www.pharmafocuseurope.com). In the context of pharma manufacturing, a digital twin might encompass: a bioreactor (with its pharmacokinetic model), a chromatography unit, a lyophilizer, an assembly robot, or a full sterile filling line. What makes it "twin" is that it continuously ingests data from its real counterpart and updates its state. Thus it can simulate future behavior (predicting, say, yield or failure modes) and allow what-if scenario testing without risking the real system.

In biopharma, digital twins have especially valuable applications:

- **Process Optimization:** By tuning process parameters (temperature, pH, agitation rate) in the virtual model, engineers can predict how to maximize output or purity, then apply those changes in the real plant.
- **Advanced Control (PAT integration):** Twin environments can integrate Process Analytical Technology sensors (e.g. near-infrared probes, Raman spectroscopy) so that both digital and physical systems share the same control algorithms. Studies report that digital twins "facilitate advanced process control productivity gains" in biologics manufacturing (www.europeanpharmaceuticalreview.com).

- **Fault Detection and Maintenance:** AI-driven twins (like the Cambridge/A*STAR model) can flag anomalies (e.g. a pump miscalibration) before batch failure, enabling predictive maintenance (www.europeanpharmaceuticalreview.com).
- **Regulatory Readiness:** Simulation and validation evidence from a digital twin must be distinguished from batch production and control records documenting actual manufacture. FDA describes records of each significant manufacturing step and identification of the people performing, supervising or checking it; simulation history alone does not establish that those steps were executed (FDA guidance).
- **Design and Commissioning:** Before a line is built or modified, the twin allows engineers to reconfigure layouts, simulate new equipment introductions, and resolve issues (like robot collision paths or material flow optimization) virtually (blogs.nvidia.com). A named example comes from automotive manufacturing: Siemens reports that Wipro PARI reduced on-site commissioning time by 70% on an engine-assembly-line project. This measures commissioning for that project, not total pharmaceutical engineering and validation time (Siemens case study).

These benefits drive strong growth: analysts project the digital twin market in pharma manufacturing to grow at ~31% CAGR through 2034 (www.worldpharmatoday.com). McKinsey describes potential gains in quality, speed, agility and resilience from digital and analytics advances in biopharma manufacturing (McKinsey analysis).

3.2 NVIDIA Omniverse and Digital Factory Twins

NVIDIA's Omniverse Enterprise is a commercial platform for building digital twins. It provides:

- A **universal 3D scene graph (USD)** that unifies data from CAD packages, BIM (building information models), and IoT sensors into one model (blogs.nvidia.com).
- Physics simulation engines for fluids, rigid bodies and thermal dynamics (important for bioreactors, HVAC, etc.).
- Libraries for robotics simulation (NVIDIA Isaac Sim) and path planning.
- AI blueprints for common tasks (e.g. video analytics, path optimization via cuOpt, etc. also relevant to factory workflow).
- Real-time ray-tracing and visual collaboration, so that engineers worldwide can view/present the factory model interactively.

For major manufacturing companies, Omniverse is being used to create *factory digital twins*. For example, Foxconn (electronics contract manufacturer) built a "Fii Digital Twin" platform with Omniverse and Siemens technology to design its factories. Foxconn can simulate entire production lines of robots and conveyors, test layouts with virtual materials handling, and even run thermal simulations of server "POD rooms" 150× faster by connecting to Omniverse-based physics solvers (blogs.nvidia.com). Similarly, TSMC (semiconductor giant) uses Omniverse to turn 2D CAD drawings of a new chip fab into an interactive 3D layout. Engineers can then virtually spot equipment collisions, plan intricate pipe routing (using NVIDIA cuOpt to auto-route multi-level piping in seconds (blogs.nvidia.com)), and evaluate how the building design impacts operations (blogs.nvidia.com). Wistron (electronics assembly) simulates robotic arms and assembly lines to improve throughput, even reducing one robot's cycle time by 12 seconds through iterative digital trials (blogs.nvidia.com).

In pharma, the use of Omniverse is still emerging but accelerating. Industry case studies (beyond Roche) include:

- **Siemens and digital pharma:** Siemens has published about using twins from lab R&D to manufacturing, pointing out that Pharma 4.0 (the intelligent factory) uses twins to realize “end-to-end digitalization” (www.pharmafocuseurope.com).
- **Sanofi’s digital labs:** While not specifically admitting Omniverse, Sanofi’s description of “digital labs” with robotic platforms implies a vision where experiments are simulated and orchestrated by AI (www.sanofi.com).
- **Academic projects:** University research, like the CARES/A*STAR twins (www.europeanpharmaceuticalreview.com), shows that tools are coming to market which could integrate with platforms like Omniverse for the commercial rollout of pharma factory twins.

Crucially, Omniverse supports **collaboration**: different teams (civil engineers, automation specialists, chemists) can work simultaneously on the same digital model. For Roche’s NC plant, NVIDIA reports that digital twins are helping accelerate facility development (NVIDIA announcement).

NVIDIA also offers a specific **AI Factory blueprint** on Omniverse for data centers themselves (pradeepstechpoints.wordpress.com). At the GTC Washington, D.C., keynote on October 28, 2025, NVIDIA founder and CEO Jensen Huang introduced the Omniverse DSX Blueprint for designing and operating gigawatt-scale AI factories (NVIDIA announcement). While Roche’s use case is different (drug & factory vs. data center design), the underlying idea is the same: unify CAD (building plans), electrical models, cooling models, and AI management into one design. In fact, Roche’s AI factory itself could be partially designed via Omniverse, since its internal supercomputing rooms and data flows need engineering. (However, presumably Roche’s data center expansion will also use similar planning tools internally.)

The bottom line: **NVIDIA Omniverse and its associated digital twin libraries are now battle-tested in high-tech manufacturing**. Roche’s innovation is applying this to pharmaceutical biomanufacturing – an area with tighter regulatory and bioprocess concerns. But the physics of fluid flow, mixing, and sterile conditions are not fundamentally different from high-tech cleanrooms. By using Omniverse, Roche aims to bring that tech-sector efficiency “off the shelf” into pharma. As Roche’s chief digital and technology officer, Wafaa Mamilli, summarized: creating digital twins for the new NC facility (and possibly others) can help “accelerate the development” of these plants (blogs.nvidia.com). NVIDIA describes digital twins as helping accelerate facility development.

4. Case Studies and Industry Context

To contextualize Roche’s efforts, it helps to compare with what other organizations are doing in AI-driven drug discovery and manufacturing:

- **Eli Lilly (USA):** On January 12, 2026, NVIDIA and Lilly announced a co-innovation lab, with the companies jointly committing up to $1 billion over five years for talent, infrastructure and compute (NVIDIA announcement). Lilly separately describes LillyPod as an NVIDIA DGX SuperPOD B300

system with 1,016 Blackwell GPUs (Lilly). In the Roche collaboration, the statement about building more sophisticated predictive models and shortening the path from biological insight to medicine came from Aviv Regev, head of Genentech Research and Early Development (Roche announcement).

- **Novo Nordisk (Denmark):** On June 11, 2025, NVIDIA announced a collaboration supporting Novo Nordisk's use of the Gefion AI supercomputer. DCAI owns and operates Gefion, which provides Novo Nordisk an AI factory for drug-discovery and agentic-AI workloads. The announcement also names BioNeMo, NeMo and Omniverse work (NVIDIA announcement). This is access to DCAI's infrastructure; Roche's announcement describes its combined on-premises and cloud footprint (Roche announcement).
- **Amgen (USA):** Amgen is also building obesity drug capacity. Its GLP-1/GIP molecule (called "MariTide") is in Phase 3, and Amgen is constructing two fill/finish plants in Holly Springs (initially $550M, then $1B added for expansion) (www.biospace.com). Amgen has publicly used AI in biologics (e.g. advanced analytics in cell line development), though not as flashy as an "AI factory." Still, the bundling of facilities by Amgen and Roche in the same region (Holly Springs) indicates intense competition for GLP-1 capacity. No specific reports state Amgen is using Omniverse, but given Amgen's size, it may also leverage digital tools (Amgen has partnership with NVIDIA in the past for GPU-accelerated drug discovery).
- **Delta Electronics, Foxconn, TSMC and Wistron (Electronics):** While not pharma companies, these manufacturers were highlighted by NVIDIA at COMPUTEX in Taipei in May 2025 for their use of digital twins in factory digitalization (blogs.nvidia.com). They serve as proof-of-concept that large-scale, high-complexity factories can be virtualized successfully. Foxconn's experience (150× faster thermal sims, plug-and-play robot fleets) and TSMC's (multi-level piping auto-routing) offer lessons that pharma can adapt. Roche's team has likely studied these examples. For instance, Omniverse's integration with Siemens and Autodesk tools (common across industries) means Roche can reuse industry-standard design data in its twins.
- **Research Collaborations:** Roche's move follows a research collaboration started in 2023 between Genentech and NVIDIA aimed at pushing generative AI for drug discovery (www.techtarget.com). Similar partnerships exist in pharma (e.g. GSK with AI startups, Pfizer with Insilico Medicine) but Roche's is distinguished by scale (the equipment deployment is bigger than typical lab-scale partnerships).

Ultimately, Roche's AI Factory initiative is both unique in scale for pharma and aligned with a broader industry shift. In 2025–2026, every major pharma/biotech is scouting ways to use AI to cut R&D costs and times. This has become a competitive necessity: even legacy companies like Novartis and smaller players have launched AI centers. Roche's approach – treating all accumulated compute as a unified utility across divisions – is ambitious. It moves AI from pilot projects to an enterprise capability.

5. Data Analysis & Evidence

This section quantifies and analyzes the expected impact of Roche's AI factory and digital twin strategies, based on available data and analogies.

5.1 AI Compute and Drug Discovery Speed

GPU compute power directly correlates with the speed and scale of AI model training. In drug discovery, larger models (with billions of parameters) can capture more chemical and biological complexity (developer.nvidia.com). By multiplying its GPU count, Roche can feasibly train much larger models than before – potentially neural networks that embed entire protein families or simulate molecular dynamics faster. While Roche has not disclosed run-times, we can infer some potential effects:

- **De novo molecule generation:** Suppose Roche trains a generative model on known inhibitors for a target. A larger model (enabled by 3,500 GPUs) can be fine-tuned faster and generate more candidate compounds in parallel. In practice, companies report that AI has designed active lead analogs in weeks rather than years. Roche cited that a backup oncology compound was discovered in **7 months** with AI, versus “more than two years” without (blogs.nvidia.com). The reported comparison is seven months versus more than two years for that program.
- **Screening and hypothesis testing:** AI can prioritize which wet-lab experiments to do. Many pharma companies describe “closing the loop” with AI to test dozens of virtual hypotheses before a single compound is synthesized. There are no published stats from Roche, but it is known in the industry that AI-driven screening can filter out 90–99% of unpromising candidates, focusing FTEs on a small set.
- **Cost per model:** GPUs also reduce cost per model: training a large model on Google TPUs or NVIDIA A100/Blackwell is much cheaper (on a per-iteration basis) than CPU or smaller GPUs. It’s beyond this report’s scope to compute ROI, but Dell/IDC reports show that NVIDIA GPUs deliver the highest performance per watt or per dollar for modern AI. Roche’s 3,500 GPUs thus represent a huge capital expenditure (costing hundreds of millions), but if they shave years off a blockbuster drug development, the ROI could be orders of magnitude in revenue (a successful new obesity drug can make ~$10–15 billion per year for its maker).

5.2 Manufacturing Efficiency Gains

Published manufacturing examples should be assessed by industry, endpoint and baseline. In Siemens’ automotive case study, Wipro PARI completed virtual commissioning of a brownfield engine assembly-line expansion in three months, compared with an estimated six to eight months under normal conditions (Siemens case study). This is a project-specific comparison, not a forecast for Roche’s North Carolina plant; it does not establish a capacity uplift, batch lead-time reduction or earlier opening for that facility.

The following evidence table separates a reported Roche result from a workload-specific vendor benchmark:

Example	Reported result	Interpretation
Genentech backup-molecule program	Seven months rather than more than two years (NVIDIA report)	A program-specific example, not an annual candidate-output forecast.
Parabricks secondary sequencing analysis	NVIDIA lists about ten minutes for 30x whole-human-genome data (NVIDIA documentation)	A vendor benchmark; performance depends on the workload and hardware.

These examples should be interpreted within their reported scope; they do not establish Roche-wide operational throughput or manufacturing gains.

Finally, one can consider energy and cost: 3,500 high-end GPUs draw substantial power (a single DGX with 8 GPUs may use ~10–15 kW). Even with liquid cooling, power costs are nontrivial. However, Roche's press emphasizes "sustainability" at the NC plant (www.roche.com), implying energy efficiency will be managed (and presumably Nvidia's current-gen GPUs are more power-efficient per flop than predecessors). Over time, as older GPUs are retired, even higher efficiency is likely.

6. Discussion and Future Directions

Roche's creation of an enterprise-scale AI factory marks a significant shift in pharmaceutical R&D strategy. Historically, drug discovery and clinical manufacture have been conservative sectors; this move signals that Roche views AI as central to its future competitiveness. Several implications and challenges include:

- **Organizational Change:** Embedding AI "across the value chain" requires retraining staff and reorganizing workflows. Many cheminformaticians, biologists, and process engineers will need to collaborate with data scientists. Organizations like Roche may need new roles (e.g. "MLops for biotech", digital twin engineers) and governance (to ensure models are validated and safe). The press release notes Roche is empowering its workforce through this tech (www.biospace.com), but real success will depend on cultural adoption.
- **Data and Validation:** High-quality data is the lifeblood of AI. Roche has vast proprietary libraries of chemical assays, genomic sequences, and clinical records – ideal for AI training. However, integrating these diverse datasets (and ensuring standardized formats) is often challenging. Digital twin models also require accurate parameters (fluid dynamics, reaction kinetics) that must be calibrated from real-world experiments. Ensuring the twin stays up-to-date (version control, calibration) will be an ongoing task.
- **Regulatory and Compliance:** Pharmaceutical regulation (FDA, EMA) currently has nascent guidance on AI/ML in drug development. If Roche uses AI models to make decisions (e.g. selecting a lead compound), it must maintain an audit trail. Similarly, any processes simulated by twins still need conventional validation. Regulators may require evidence that the AI-assisted decisions meet the same safety standards. On the manufacturing side, a digital twin's simulations may need to be qualified by real test batches. The fact that digital twins can store compliance data (as noted by Cambridge researchers (www.europeanpharmaceuticalreview.com)) is promising, but Roche will still have to demonstrate to regulators that "running the plant virtually" did not introduce any risk that wasn't observed in bench or pilot runs.
- **Cybersecurity:** With centralized AI labs and cloud integrations, ensuring data security is paramount. Healthcare and pharma are high-value targets for cyber espionage. NVIDIA's enterprise solutions and cloud partnerships emphasize encryption and compliance, but any breach could compromise medical IP or patient data. Roche will likely implement strict isolation (e.g. Genentech's DGX clusters may be in locked labs) and move only encrypted datasets through cloud.

- **Economic Impact:** The capital expenditure (CAPEX) here is enormous – 3500 GPUs plus infrastructure likely represents hundreds of millions of dollars investment. However, if it results in just one blockbuster drug or a fraction of a percent in productivity, the returns could be hundreds of millions per year. Analysts will watch industry ROI: if Roche's pipeline advances accelerate significantly, competitors may be compelled to match this scale. Already Lilly's $1B and Roche's effort suggest an "AI arms race" in pharma R&D.

Looking ahead, several future directions emerge:

- **Agentic Research:** Roche alluded to "agentic research workflows" (blogs.nvidia.com) – the idea that AI agents could autonomously design and conduct experiments in the lab. With a powerful AI factory, Roche might pursue closed-loop robotics labs where an AI model generates a hypothesis, programs a robot to synthesize a compound or run an assay, interprets the result, and iterates. NVIDIA and partners have demos of such "AI scientists." Complete digital twins of labs (beyond just factories) could allow Roche to simulate yeast fermenters or cell culture processes long before an experiment.
- **Personalized Digital Twins:** The idea of a digital twin isn't limited to hardware. In the broad vision of "Precision Medicine," one could imagine a patient-level digital twin. Combined with clinical and genomic data, Roche could one day run virtual trials on in silico patient cohorts, accelerating drug candidate selection and dosing decisions. While speculative, Roche's overlap of diagnostics (clinical data) and therapeutics makes it uniquely positioned to explore this synergy. NVIDIA has shown early work on digital health twins and generative AI for patient stratification.
- **Integration with Robotics:** NVIDIA's Isaac Sim and Jetson platforms can integrate with Omniverse to form "cyber-physical systems." For manufacturing in a new plant, floor robots (AGVs, robotic arms) could be co-developed with the plant twin. For example, a robot's path in the digital twin could be optimized before physical deployment, ensuring no unexpected collisions with equipment. Real-time sensor data from robotic end-effectors could feed back into maintaining the twin's accuracy.
- **Generative Manufacturing:** Looking further, one can imagine generative design applied to the plant itself. Omniverse's reinforcement learning (via Isaac Lab and cuOpt) could automatically propose an optimal floorplan or piping layout that humans might not conceive. Already in electronics, companies use NVIDIA's copper-routing AI to layout PCB traces. In pharma, a similar algorithm could propose novel bioreactor shapes or mixing strategies. Roche's investment signals that such advanced tools will be tested in biomanufacturing.

7. Conclusion

Roche's announcement of a **3,500+ GPU AI factory with NVIDIA** and the deployment of Omniverse digital twins for an upcoming GLP-1 facility represents a watershed moment in pharmaceutical innovation. In March 2026, Roche described its hybrid-cloud AI factory as the pharmaceutical industry's largest announced, with more than 3,500 GPUs (Roche announcement). By building this AI infrastructure, Roche aims to do "now what patients need next" (blogs.nvidia.com): get safe, effective medicines and diagnostics into hands faster and more efficiently.

The immediate goal is clear: accelerate Roche's discovery programs (using *BioNeMo*-trained models), make clinical development more data-driven (advanced imaging/genomics analysis), and build plant designs faster and smarter (Omniverse twins for the Holly Springs GLP-1 plant and others). In doing so, Roche will both bolster its position in the booming obesity/GLP-1 market and lay the foundation for a decade of AI-driven productivity.

From a broader perspective, Roche's AI factory exemplifies a new paradigm in pharma. Too long has drug development relied on manual processes and siloed tools. Now, with cluster-scale AI, companies can integrate data and operations end-to-end. The implications extend beyond GLP-1 drugs: any complex biologic (antibodies, cell therapies, vaccines) could be designed or produced more quickly. Moreover, the diagnostics side – traditionally separate – can be fully merged with therapeutics (the so-called "theranostics" model), potentially leading to personalized treatment plans co-developed by the AI.

However, realizing this promise will require careful execution: ensuring regulatory readiness, maintaining data integrity, and measuring actual patient outcome gains. Stakeholders will expect evidence that AI actually improves workflows without introducing new risks. Transparency of models and validation against experiments will be key.

If successful, Roche's approach may become the blueprint for others. Already, NVIDIA and other tech providers are promoting "AI Factory" blueprints to life sciences. Within a few years, one might see similar announcements from Novartis, Johnson & Johnson, or even smaller biotech consortia. The question then becomes not whether AI will transform pharma, but how fast and in what new directions. Roche's AI factory – with its 3,500 GPUs and digital twins – is propelling that transformation forward.

References: This report synthesized information from Roche and NVIDIA press releases, industry news, and research reports. Key sources include Roche's official media releases (www.biospace.com) (www.roche.com), the NVIDIA/Genentech blog (blogs.nvidia.com), technology news articles (www.techtarget.com) (www.datacenterdynamics.com), and market analyses (www.jpmorgan.com) (www.globenewswire.com) (www.worldpharmatoday.com). Citations are provided inline in [brackets] linking to the reference statements.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.roche.com/media/releases/med-cor-2026-03-16
2. https://www.biospace.com/press-releases/roche-launches-nvidia-ai-factory-to-accelerate-the-development-of-new-therapeutics-and-diagnostics-solutions
3. https://www.datacenterdynamics.com/en/news/pharmaceutical-company-roche-deploys-3500-nvidia-blackwell-gpus-across-hybrid-cloud-and-onpremises
4. https://blogs.nvidia.com/blog/roche-ai-factories-omniverse
5. https://www.globenewswire.com/news-release/2026/01/07/3214628/28124/en/GLP-1-Market-Trends-and-Global-Forecasts-Report-2025-A-157-5-Bn-Market-by-2035-Landscape-Led-by-China-and-the-U-S-Rising-Partnerships-and-Strong-DTC-Driven-Demand-Push-Biologics-Ab.html
6. https://www.jpmorgan.com/insights/global-research/current-events/obesity-drugs
7. https://www.roche.com/media/releases/med-cor-2025-08-25

8. https://www.biospace.com/business/how-north-carolina-attracted-amgen-and-roche-to-become-a-next-gen-obesity-drug-production-hotspot
9. https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-the-future-of-biopharma-manufacturing
10. https://www.techtarget.com/pharmalifesciences/news/366640421/Roches-new-NVIDIA-deal-locks-in-the-industrys-largest-hybrid-cloud-AI-factory
11. https://www.roche.com/investors/annualreport25
12. https://www.accessdata.fda.gov/drugsatfda_docs/label/2026/218316s007lbl.pdf
13. https://www.novonordisk.com/news-and-media/news-and-ir-materials/news-details.html?id=916407
14. https://www.fiercebiotech.com/biotech/roche-inks-27b-carmot-buyout-muscle-obesity-market-securing-injectable-and-oral-assets
15. https://www.roche.com/media/releases/med-cor-2026-06-01
16. https://www.roche.com/irp260723-a.pdf
17. https://www.genentech-medinfo.com/pipeline-and-clinical-trials/cvrm/ct-996.html
18. https://www.fiercepharma.com/pharma/roche-begins-construction-7000m-obesity-drug-manufacturing-facility-north-carolina
19. https://nvidianews.nvidia.com/news/nvidia-and-lilly-announce-co-innovation-lab-to-reinvent-drug-discovery-in-the-age-of-ai
20. https://blogs.nvidia.com/blog/omniverse-digital-twins-taiwan-manufacturers-physical-ai
21. https://www.europeanpharmaceuticalreview.com/ai-digital-twin-developed-to-boost-pharma-manufacturing-efficiency/265069.article
22. https://www.pharmafocuseurope.com/articles/digital-twins-in-pharma
23. https://pradeepstechpoints.wordpress.com/category/ai-factories
24. https://developer.nvidia.com/blog?p=104887
25. https://docs.nvidia.com/clara/latest/Overview.html
26. https://docs.nvidia.com/nemo/guardrails/configure-guardrails/guardrail-catalog/fact-checking.html
27. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cpg-sec-425500-computerized-drug-processing-identification-persons-batch-production-and-control
28. https://resources.sw.siemens.com/en-US/case-study-wipro-pari
29. https://www.worldpharmatoday.com/biopharma/using-digital-twins-to-optimize-pharmaceutical-plant-performance
30. https://www.sanofi.com/en/magazine/ai-in-healthcare/ai-research-development-value-chain-manufacturing-digital-labs-self-sharpening-tools
31. https://blogs.nvidia.com/blog/omniverse-dsx-blueprint
32. https://www.lilly.com/news/stories/new-supercomputer-could-change-future-medicine
33. https://nvidianews.nvidia.com/news/nvidia-partners-with-novo-nordisk-and-dcai-to-advance-drug-discovery

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.