



NEWS RELEASE

Lantern Pharma Launches ZetaOmics™ — an Autonomous, Intelligent “Computational Biologist” That Brings Expert-Grade Bioinformatics, Biostatistics & Reasoning to Every Cancer

2026-07-08

First previewed in the company's May roadmap, the new withZeta.ai module now enters early access with leading bioinformatics teams and Lantern collaborators — autonomously designing experiments, building rigorous cohorts from 615,000+ curated cancer samples, and delivering publication-ready multi-omic analysis in minutes

- A new class of research agent. ZetaOmics™ is an autonomous “Computational Biologist” that goes beyond automating software — it reasons about experimental design, prevents pitfalls that produce false discoveries, and makes statistically rigorous decisions on its own, encoding the judgment of a tenured bioinformatician so that scientific rigor is the default across rare and common cancers alike.
- Intelligent discovery, analysis, and cohort-building at scale. Through direct access to more than 615,000 curated cancer samples across eleven data sources, ZetaOmics resolves fragmented data and vocabularies autonomously and builds meaningful, analysis-ready cohorts from natural-language questions — turning weeks of manual data wrangling into minutes.
- Publication quality output that accelerates insights. From differential expression and pathway enrichment to survival modeling, drug-response and protein-level validation, ZetaOmics delivers auditable, publication-quality results and figures inside a single conversational workflow — changing and accelerating how bioinformatics and computational biology get done.

DALLAS--(BUSINESS WIRE)-- Lantern Pharma Inc. (NASDAQ: LTRN), a clinical-stage, AI-driven precision oncology company, today announced the launch of **ZetaOmics**, the computational-biology module of its multi-agent AI co-scientist platform, **withZeta.ai**. ZetaOmics introduces an autonomous “Computational Biologist” persona (agent) that performs real, end-to-end bioinformatics and multi-omic analysis — across any cancer type, and is purpose-built for the rare and pediatric cancers that have long been underserved by dedicated computational resources and meaningful bioinformatics.

First introduced as part of the withZeta.ai development roadmap the Company unveiled on May 7, 2026, ZetaOmics now moves from roadmap to reality. It is being made available initially through an **early-access program** for a select group of leading academic and industry bioinformatics teams and Lantern Pharma collaborators, whose real-world use will guide refinement ahead of broader commercial availability. This phased rollout is designed to validate the module against the most demanding research workflows while establishing the reference relationships that seed a larger subscriber and partnership base.

The current generation of AI research tools forces a choice between a brilliant generalist that has read everything but can run nothing, and legacy pipelines that execute fixed routines but cannot reason. ZetaOmics is built as a third option: an agent that both reasons and runs — designing the analysis, executing it independently on real biological data, defending its methodological choices, and returning publication quality results to provide true real-time value to researchers.

Intelligence embedded in every tool & module — not just automation

Most emerging agentic bioinformatics systems wrap a language model around off-the-shelf software, automating command-line execution. That approach can still build a confounded cohort, choose the wrong statistical test, or compare datasets that were never comparable — returning a confident answer that only an expert would recognize as wrong. ZetaOmics takes the opposite approach: it operates on pre-computed, harmonized, multi-omic data layers with domain intelligence embedded directly into each of its fourteen tools. Where conventional systems will run a flawed analysis and hand back flawed results, ZetaOmics is designed to recognize the flaw, decline to run it, and explain how to fix the experimental design.

Every analysis executes on withZeta.ai’s production runtime with unified authentication and per-execution logging — who asked what, which tools ran, which data was accessed — producing a queryable, exportable audit trail suited to regulated, compliance-sensitive research.

A Defining Step for AI in Oncology

“For decades, the deepest bottleneck in drug discovery hasn’t been data; it has been judgment. The rare instinct of a great bioinformatician or computational cancer biologist—to know which test the data calls for, when two datasets should not be compared, and when a result is true signal rather than statistical noise—is scarce, costly, and often lost when that expert leaves. With ZetaOmics, we’ve worked to encode that judgment into an autonomous agent and make rigor the default — so that any researcher, anywhere, can run analysis at a level that was once reserved for the best-resourced labs in the world.

We believe this is what the next decade of oncology R&D will be built on: co-scientists that don’t just retrieve knowledge, but generate it — responsibly, reproducibly, and at a pace that finally matches the urgency facing patients. By opening ZetaOmics first to the world’s leading bioinformatics teams and our collaborators, we intend to prove its value where the science is hardest, and in doing so open entirely new markets, new partnership and collaboration opportunities, and new subscription revenue for Lantern — across every cancer, not only the rare ones. This is how we scale withZeta.ai into a durable, non-dilutive growth engine for our shareholders while also advancing our own pipeline,” said Panna Sharma, President and Chief Executive Officer of Lantern Pharma and Founder of withZeta.ai.

Judgment and automation in action

The difference between orchestrating tools and encoding judgment is easiest to see in practice. In each example below, a single natural-language request triggers a chain of expert decisions that a skilled bioinformatician would make by hand over days or weeks — executed autonomously, with the reasoning made explicit and auditable.

1 · Finding a drug vulnerability in a rare, hard-to-treat cancer

The ask: “Are there drug vulnerabilities in KRAS-mutant lung cancer, and are they real at the protein level?”

What ZetaOmics does autonomously:

- Builds batch-safe cohorts, checking for mixed pipelines and confounder risk before any comparison runs.
- Enforces functional consequence, so only true KRAS loss-of-function variants define the mutant group and benign missense noise doesn’t dilute the signal.
- Checks data coverage and power, confirming enough mutant and wild-type lines carry drug and CRISPR data, and reporting effect sizes, confidence intervals, and FDR — not just p-values.
- Validates hits at the protein level via RNA-protein concordance, distinguishing a druggable membrane target from a transcript with no protein behind it.

The impact: A vulnerability that a specialist team might spend weeks qualifying — and could easily over-claim —

is surfaced, statistically defended, and protein-validated in a single conversational session.

2 · Refusing a comparison that would have produced a false discovery

The ask: “Compare gene expression between these two glioblastoma cohorts I’ve assembled.”

What ZetaOmics does autonomously:

- Profiles the cohorts first, detecting that they span multiple independent projects and mix incompatible quantification pipelines (RSEM vs. Kallisto).
- Declines to run the naive comparison, because the batch difference would dominate the biology and generate a confident but false result.
- Explains the confounder and fixes the design, auto-matching pipeline-consistent, anatomically correct controls so a valid contrast can run.

The impact: The single largest source of false discoveries in multi-cohort studies is caught and corrected automatically — the kind of quiet, judgment-based save that normally depends on a senior analyst noticing.

3 · Prioritizing a druggable biomarker in a pediatric cancer with sparse data

The ask: “What are the most promising druggable targets in this rare pediatric tumor?”

What ZetaOmics does autonomously:

- Assembles a multi-evidence target shortlist— unfavorable-prognostic, tumor-restricted, secreted or membrane-accessible — checked against normal-tissue baselines.
- Recognizes the data gap: with no validated prognostic panel available for this rare indication, it automatically pivots to a plasma-led path.
- Surfaces circulating biomarker candidates— proteins measurably elevated in patient plasma versus healthy controls — as the actionable shortlist.

The impact: For exactly the populations the field tends to abandon, the agent adapts its own method to the available evidence rather than returning nothing — turning data scarcity into a defensible starting point.

Inside the Computational Biologist: 14 intelligent tools

ZetaOmics equips the Computational Biologist persona (agent) with fourteen prioritized tools. Each carries embedded domain intelligence that makes rigorous choices automatically, rather than leaving them entirely to the

researcher:

MODULE	WHAT SETS IT APART	AUTONOMOUS CAPABILITY
Discovery	Ask in plain English	Federated search across 615K+ samples and 11 sources; auto-resolves cancer-type naming and routes only to datasets that hold the modality you need.
Gene Intelligence	Genome to proteome in one query	Links gene, transcript, and protein across Ensembl, UniProt, OMIM, and Orphanet, with pre-computed co-expression networks and rare-disease associations.
Batch Guardian	Stops false discoveries before they start	Profiles cohorts for hidden confounders and mixed pipelines before any analysis runs, and auto-matches anatomically correct normal controls.
Cohort Builder	Describe it; get an analysis-ready group or cluster	Turns plain-language filters into sample groups with diagnostic feedback — if a filter finds nothing, it returns the valid values that do exist.
Expression	The right normalization, chosen for you	Auto-selects normalization by task, scores curated biological signatures, resolves canonical pathways, and validates RNA against protein evidence.
Mutations	Not all variants are equal	Consequence-aware stratification (true loss-of-function vs. benign missense) and copy-number analysis with druggability context; live cBioPortal access.
Drug Response	The full statistical story	Resolves drug synonyms across ChEMBL/PubChem, checks data coverage before comparing, and reports effect size, confidence intervals, power, and FDR — not just p-values.
Differential Expression	Rigor enforced automatically	Adaptively selects edgeR, DESeq2, or limma-trend by sample size, refuses cross-pipeline and cross-specimen comparisons by default, and returns pathway enrichment with caveats surfaced.
Survival	Beyond a median split	Covariate-adjusted Cox modeling with hazard ratios, confidence intervals, and proportional-hazards diagnostics — flagging unadjusted associations as potentially confounded.
Stratification	Test any molecular hypothesis	Splits cohorts by expression, gene ratios, curated signatures, clinical metadata, or drug/CRISPR response into analysis-ready groups for downstream testing.
Protein Triage	Is the target actually druggable?	Ranks druggable, tumor-restricted targets against normal-tissue baselines — and for rare cancers with no prognostic panel, pivots automatically to plasma-based biomarker discovery.
Analysis Sandbox	Bespoke tests, no code required	Describe a custom analysis in plain language; the agent generates and runs sandboxed code against stored results, routing to guarded purpose-built tools whenever they apply.
Visualization	Publication-quality, every time	Thirteen presets (volcano, Kaplan-Meier, drug waterfall, dependency bars, concordance, PCA/UMAP and others) rendered deterministically so every figure matches the reported numbers.
Investigation Memory	Never rebuild from scratch	Stores, versions, and traces every result across turns, so complex multi-step workflows checkpoint their work and stack more analysis than any single run could survive.

Together, these tools let a researcher move from a single natural-language question to a rigorous, protein-validated, publication-ready answer — with no infrastructure to manage, no pipelines to debug, and batch-effect safeguards enforced automatically.

Built for the economics and urgency of rare cancer

Rare cancers account for roughly a fifth of all cancer cases, yet development economics rarely justify dedicated computational-biology teams for populations of fewer than 200,000 patients — a systematic gap ZetaOmics is designed to close. By making expert-grade analysis available on demand across all cancer types, ZetaOmics extends withZeta.ai's reach across the full oncology landscape while remaining purpose-built for the rare and pediatric settings where analytical expertise is scarcest.

Opening new markets, partnerships, and recurring revenue

For Lantern Pharma shareholders, ZetaOmics is designed to expand the commercial surface area of withZeta.ai well beyond its rare-cancer origins. Because the Computational Biologist operates across all cancer types — and because rigorous multi-omic analysis is a daily need for essentially every oncology research group — the module opens addressable demand across academic medical centers, biotech and pharmaceutical R&D teams, integrated cancer centers, and contract research organizations that were not previously core to the platform.

The Company expects ZetaOmics to create three distinct avenues for value creation:

- New markets. Extending expert-grade computational biology to all cancer types broadens withZeta.ai's reach across the full oncology landscape, positioning the platform to capture demand in a market for AI-enabled drug-discovery tools that independent analysts size in the multiple billions of dollars and project to grow at double-digit annual rates through the early 2030s.
- New collaboration and partnership opportunities. The early-access program is structured to convert leading bioinformatics teams and institutional collaborators into reference partners — seeding co-development relationships, enterprise deployments, and data and technology partnerships with cancer centers, research organizations, and large biopharma.
- New subscription potential. As a natively subscription-based capability layered onto withZeta.ai's existing tiered model, ZetaOmics is intended to deepen engagement and expand recurring, non-dilutive revenue — supporting higher-value commercial and enterprise tiers and creating natural upsell paths as teams move from individual analysis to organization-wide deployment.

Together, these dynamics are intended to strengthen withZeta.ai's position as both a hyper-productivity engine for Lantern's own drug-discovery pipeline and a scalable, standalone commercial platform for the global oncology research community.

About withZeta.ai™

withZeta.ai is redefining how rare cancer research, discovery, drug development, and clinical trial design gets done. Knowledge work in oncology is migrating to AI co-scientists, autonomous systems that investigate, reason, and synthesize across the full breadth of scientific evidence. **withZeta.ai** is that co-scientist: purpose-built for the biology, economics, and urgency of rare cancer drug development, and accessible to any researcher, anywhere. Built by Lantern Pharma (Nasdaq: LTRN). Learn more and subscribe at **withZeta.ai**.

About Lantern Pharma

Lantern Pharma (NASDAQ: LTRN) is a clinical-stage AI-driven precision oncology company transforming the cost, pace, and timeline of oncology drug discovery and development. The company's proprietary AI and machine

learning platform, RADR®, now operationalized through withZeta.ai, leverages billions of data points and advanced computational methods to rapidly uncover biomarker signatures and accelerate the development of targeted oncology therapies for difficult-to-treat cancers, including those of the central nervous system. Lantern is currently advancing a pipeline of small molecule drug candidates and antibody-drug conjugate programs across multiple solid tumor and hematologic malignancies. For more information, visit www.lanternpharma.com.

Forward-Looking Statements

This press release contains forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended. These forward-looking statements include, among other things, statements relating to: future events or our future financial performance; the planned development, design, capabilities, and commercial potential of the withZeta.ai platform and the ZetaSwarm™ and ZetaOmics™ feature tracks; expectations about enterprise grade platform capabilities and anticipated implementation and deployments; our broader product, clinical, and platform roadmap, including future directions for withZeta.ai and the planned implementation of new features; the potential advantages of our withZeta.ai platform in identifying drug candidates, accelerating drug development, and generating revenue through software licensing and subscription models; the planned commercialization of our AI platforms, including withZeta.ai, and the expected market opportunity for AI co-scientist platforms; our intention to leverage artificial intelligence, machine learning and genomic data to streamline and transform the pace, risk and cost of oncology drug discovery and development; and estimates regarding potential markets and potential market sizes.

Any statements that are not statements of historical fact (including, without limitation, statements that use words such as “anticipate,” “believe,” “contemplate,” “could,” “estimate,” “expect,” “intend,” “seek,” “may,” “might,” “plan,” “potential,” “predict,” “project,” “target,” “model,” “objective,” “aim,” “upcoming,” “should,” “will,” “would,” or the negative of these words or other similar expressions) should be considered forward-looking statements. There are a number of important factors that could cause our actual results to differ materially from those indicated by the forward-looking statements, such as (i) the risk that we may not be able to secure sufficient future funding when needed and as required to advance and support our existing and planned development programs and operations, (ii) the risk that our AI platform commercialization efforts, including withZeta.ai, may not generate the anticipated revenue or achieve the expected market adoption, (iii) the risk that implementation of our product development plans and new features for withZeta.ai may not be successful or may take longer than anticipated for development, completion, and release, (iv) the risk that no drug product based on our proprietary AI platforms has received FDA marketing approval or otherwise been incorporated into a commercial product, and (v) technical, scientific, regulatory, financial, competitive, and operational risks and those other factors set forth in the Risk Factors section in our Annual Report on Form 10-K for the year ended December 31, 2025, filed with the Securities and Exchange Commission on March 30, 2026.

You may access our Annual Report on Form 10-K for the year ended December 31, 2025 under the investor SEC filings tab of our website at www.lanternpharma.com or on the SEC's website at www.sec.gov. Given these risks and uncertainties, we can give no assurances that our forward-looking statements will prove to be accurate, or that any other results or events projected or contemplated by our forward-looking statements will in fact occur, and we caution investors not to place undue reliance on these statements. All forward-looking statements in this press release represent our judgment as of the date hereof, and, except as otherwise required by law, we disclaim any obligation to update any forward-looking statements to conform the statement to actual results or changes in our expectations.

Investor & Media Contact

Investor Relations: ir@lanternpharma.com | +1-972-277-1136

Source: Lantern Pharma Inc.