



Concentriq Compute

Get your AI models from ready
to running in days.

A fully managed compute service
for running your pathology
AI where your data lives — on
Concentriq. Deploy and scale
without IT procurement or
management cycles.

Why Production-Ready Models Rarely Reach Production

Pathology AI is changing how life sciences organizations discover, develop, and deliver therapies. Biomarker discovery programs are identifying predictive signatures across thousands of whole slide images. Clinical trial teams are using AI to surface complex tissue signatures that sharpen cohort selection and increase the odds of a meaningful readout.

The pressure to build and run these models at scale is growing, and data scientists and developers have answered the call. With AI coding assistants and open-source foundation models, building production-ready models is faster and more accessible than ever. Yet most of their models never actually make it to production.

Standing up a compute environment for pathology data means months of cloud provisioning, compliance validation, and security documentation — all before a

single image is processed. Research teams can't execute this independently, and IT teams absorb these requests into backlogs that move on quarterly cycles, not research timelines. Every new algorithm restarts the process, and the overhead compounds.

Even when teams overcome these barriers, running data outside its governed environment makes maintaining security and compliance difficult. And in many cases, close to impossible.

Meanwhile, the models are losing value while sitting in limbo. Study readouts arrive late, trial windows narrow, and the evidence needed to inform the next decision isn't generated in time. Every quarter a model sits idle is a quarter of unrealized return on the R&D investment that built it.

Concentriq Compute eliminates the infrastructure and operational complexity standing between your models and production. It's a fully managed compute service built on Concentriq, Proscia's enterprise pathology platform.

You provide a containerized workload. Concentriq handles scheduling, execution, scaling, and lifecycle management. Results land directly in the platform — connected to every source image, ready for review and downstream analysis.



Ready to Running in Days

Your scientifically ready model shouldn't wait months for infrastructure to catch up. With Concentriq Compute, you move from first model run to production-scale deployment in days – without waiting on server provisioning, IT approvals, or procurement cycles that were never designed for the pace of research.

Scale Your Models, Not Your Overhead

Your fifth model deploys the same way as your first – and so does your fiftieth. Infrastructure burden stays flat no matter how many programs your team runs. Scale across therapeutic areas and study cohorts without waiting for IT capacity or budget to open up. When a model fails mid-run, recovery happens automatically. When demand spikes, resources scale to meet it.

No Data Exports, Duplication, or Transfers

Run models directly where the data lives. Sensitive pathology data stays within a single governed, GxP-ready environment. Security review cycles, validation burden, and compliance exposure associated with data movement are eliminated entirely.

Every Run Grows Your Data Foundation

AI outputs return to Concentriq connected to the source image, metadata, and analysis run that produced it. Over time, those results compound into an increasingly rich data foundation – scores, classifications, and more – all searchable and traceable to the source. Query across everything the organization has ever generated to fuel new model development, inform portfolio strategy, and accelerate the next decision.



Everything Your AI Program Needs to Scale

4



No IT Dependency

Submit models and access compute directly. A container and a configuration file is all that stands between a model and its first production run.



Automated Recovery

Failed jobs never stall a program. Recovery and lifecycle management are handled automatically, so your team isn't troubleshooting infrastructure.



Elastic Auto-Scaling

GPU, CPU, and memory resources provision on demand and release when idle. You pay for what you use — no overprovisioning, no wasted spend.



Predictable Resource Allocation

Set maximum resource consumption thresholds to keep costs within budget while still enabling auto-scaling.



Security and Compliance, Built In

Your models run on the same enterprise-grade security infrastructure your data already lives on. Certified for compliance with ISO 27001 and SOC 2 Type II, designed to support medical devices under ISO 13485, and when specifically configured, HIPAA- and GxP-compliant.



Secured IP

Your proprietary models and source code stay yours — secured in dedicated, isolated compute environments with end-to-end encryption at rest and in transit.



Traceable, Reproducible, Auditable

Every structured quantitative result (including provenance metadata and file attachments) is persistently linked to the source image, the pipeline that produced it, and the execution timestamp.



Searchable AI Data at Scale


As programs accumulate, Concentriq becomes a queryable foundation across every AI result the organization has produced. Search across studies, cohorts, and therapeutic areas to surface patterns.



Up and Running in Days

Getting your model into production takes three inputs: a containerized workload, your CPU or GPU requirements, and a preference for triggering runs through the Concentriq interface or API. From there, Proscia manages the rest.

Proscia can also provide CI tooling to help teams build container images, so packaging your model is supported from the start.



Your Models, Running Where Pathologists Already Work


Models integrated through Concentriq's Image Analysis API appear directly in the workflows pathologists and scientists use every day. Trigger analysis through the Concentriq viewer as part of routine image review, or programmatically via API for batch processing and automated pipelines.

AI outputs return as overlays, heatmaps, and structured results within the Concentriq viewer, alongside every other image analysis tool already running on the platform. One system of work, one familiar experience.

This is what Concentriq Compute makes possible: your proprietary and open-source models, running on fully managed infrastructure, delivering results directly into the platform your teams already trust. No infrastructure to build or workflows to change. Just models in production, generating value.

EXAMPLE

A team has fine-tuned a cell detection model in QuPath — one of digital pathology's most widely adopted open-source analysis platforms — and needs to run it at production scale. With Concentriq Compute and Image Analysis API, deploying that model requires only a processor container and a configuration file. The result is a full production deployment delivered in days, with scientists viewing cell detection results rendered directly as overlays alongside the original whole slide image in Concentriq. The next model deploys the same way.




Your AI — Built and Deployed at Enterprise Scale

7

Concentriq Compute requires only Concentriq — no additional Proscia products. But for organizations looking to fuel and build their models too, it pairs with Concentriq Embeddings and multimodal real-world data for a development-to-deployment environment within a single platform.

However you build your AI, Proscia is invested in helping it reach production and deliver results.



MULTIMODAL REAL-WORLD DATA

Train and validate AI models.

Access whole slide images linked to clinical, molecular, and genomic context to augment controlled clinical trial data and support the development of high-performing, clinically viable AI models.



CONCENTRIQ EMBEDDINGS

Accelerate model building.

Generate embeddings from state-of-the-art foundation models directly in Concentriq, expediting AI model development 13x faster than traditional approaches.



CONCENTRIQ COMPUTE

Deploy and scale instantly.

Get your proprietary AI models from ready to running in days — directly where your data lives. Retain full IP ownership and scale from pilot workloads to large studies — without provisioning servers or managing infrastructure.



Get Started

Discover how Concentriq Compute can accelerate your AI programs.

Connect with your Proscia account team to schedule a personalized demo.

[PROSCIA.COM/CONTACT](https://proscia.com/contact)



**2026
BEST IN KLAS**

**RATED #1
BY CUSTOMERS**

© Copyright 2026. Proscia Inc. Proscia, Concentriq, Proscia Aperture, the circles logo and Proscia Aperture logo are trademarks or registered trademarks owned by Proscia Inc. All rights reserved.

Concentriq AP-Dx is FDA-cleared for primary diagnosis in the United States (U.S.) with the Hamamatsu NanoZoomer® S360MD Slide scanner and is licensed in Canada and CE-marked under the EU IVDR for primary diagnosis. Concentriq AP-Dx is not intended for use with frozen sections, cytology, or non-FPPE hematopathology specimens. Concentriq AP and Concentriq LS are for Research Use Only. Not for use in diagnostic procedures. Proscia's AI applications are available for research use only.