NonStop

# The AI Readiness Assessment for Genomics Platforms

A 15-minute self-assessment to know if your platform can actually support production AI before you waste months and thousands of dollars finding out the hard way.

We built this after watching three clients discover mid-implementation that their platforms couldn't support the AI capabilities they'd committed to building. Use this before you sign vendor contracts or allocate engineering resources.

## The 10-Question AI Readiness Assessment

Answer honestly. Every "no" represents technical debt you'll need to address before AI works in production.

### Data Infrastructure

**Q1: Can you retrieve the exact version of every reference database used to process a specific patient sample from 18 months ago?**

- ☐ Yes - We version ClinVar, gnomAD, dbSNP, and can reconstruct the exact state from any date
- ☐ Partial - We log database versions but can't easily reconstruct past states
- ☐ No - We use "latest" versions without specific tracking

**Why this matters:** AI training requires knowing exactly what data the model learned from. Without database versioning, you can't reproduce predictions or validate model updates.

---

**Q2: Do you track computational lineage from raw FASTQ to final variant calls?**

- ☐ Yes - Every intermediate file, tool version, and parameter is logged with sample linkage
- ☐ Partial - We log some steps, but lineage is incomplete
- ☐ No - We track inputs and outputs, but not transformation steps

**Why this matters:** AI models need features computed from your pipeline (coverage, quality scores, allele frequencies). If you can't trace how those features were calculated, you can't validate AI predictions.

---

**Q3: Can you programmatically access patient phenotype data in a structured format?**

- ☐ Yes - HPO terms or structured diagnosis codes with API access
- ☐ Partial - Phenotype data exists, but it's in free-text notes or PDFs
- ☐ No - Phenotype information isn't systematically captured

**Why this matters:** AI variant prioritization uses phenotype matching. Free-text clinical notes don't work; you need structured HPO (Human Phenotype Ontology) terms or ICD codes.

---

### Model Operations

**Q4: Do you have GPU compute infrastructure provisioned (or easy cloud access)?**

- ☐ Yes - On-demand GPU access for training and inference
- ☐ Partial - We can request GPU instances but it takes days/weeks
- ☐ No - GPU infrastructure or budget allocated

**Why this matters:** Deep learning models (DeepVariant, SpliceAI, AlphaGenome integration) require GPUs. CPU-only inference is 10-50x slower and often impractical for clinical turnaround times.

## Validation & Compliance

**Q5: Do you have a documented validation protocol for introducing new algorithms?**

- ☐ Yes - Written SOPs covering validation sample size, metrics, documentation
- ☐ Partial - We validate, but it's ad-hoc, not standardized
- ☐ No - formal validation process exists

**Why this matters:** CAP requires documented validation for any algorithm affecting clinical decisions. AI is no exception. Ad-hoc validation fails inspections.

---

**Q6: Can you generate an audit report showing every person who accessed a patient's genomic data?**

- ☐ Yes - Comprehensive audit logs with user, timestamp, action, and data accessed
- ☐ Partial - We log some access, but not comprehensively
- ☐ No - systematic audit logging

**Why this matters:** HIPAA requires this. AI systems accessing patient data need the same audit trails. This is table stakes for clinical deployment.

---

**Q7: Do you have a documented incident response plan for AI prediction errors?**

- ☐ Yes - Written procedures for detecting, investigating, and remediating AI errors
- ☐ Partial - General incident response exists but not AI-specific
- ☐ No - formal incident response procedures

**Why this matters:** AI will make wrong predictions. How you detect and respond to errors determines patient safety and regulatory compliance.

---

## Operational Readiness

**Q8: Can your genetic counselors or lab scientists override AI predictions with documented**

- ☐ Yes - On-demand GPU access for training and inference
- ☐ Partial - Overrides are possible but not systematically documented
- ☐ No - No override mechanism exists

**Why this matters:** AI assists human experts; it doesn't replace them. Expert override capability with documentation is a regulatory requirement and clinical safety necessity.

---

**Q9: Do you measure per-sample computational cost?**

- ☐ Yes - We track compute, storage, and egress costs per sample/pipeline
- ☐ Partial - We see aggregate cloud bills but not per-sample attribution
- ☐ No - We don't track computational costs at the sample level

**Why this matters:** AI increases computational costs 2-5x. Without per-sample cost tracking, you can't budget accurately or optimize effectively.

---

**Q10: Can you retrain your AI models with updated clinical evidence without breaking production?**

- ☐ Yes - Model retraining pipeline with staging, validation, versioned deployment
- ☐ Partial - We can retrain but deployment disrupts production
- ☐ No - Models are static; retraining means rebuild

**Why this matters:** ClinVar updates 2,000+ variants monthly. Static AI models become outdated. Retraining capability determines whether AI evolves or ossifies.

## Scoring Your Platform

Count your "Yes" answers

| | |
|---|---|
| • 9-10 Yes | Your platform is AI-ready. Focus on use case selection and validation studies. |
| • 6-8 Yes | You're close but have gaps. Budget 4-6 months and $200-400K to fill critical infrastructure gaps before deploying production AI. |
| • 3-5 Yes | Significant infrastructure work required. You're 9-12 months and $400-700K from an AI-ready platform. Consider whether AI ROI justifies this investment vs. other priorities. |
| • 0-2 Yes | Your platform isn't ready for production AI. Attempting deployment now will result in failed validations, inspection findings, or expensive rework. Budget 12-18 months and $600K-1M for foundational infrastructure before pursuing AI capabilities. |

## One More Thing: The Vendor Question Test

**When evaluating AI vendors, ask them these three questions:**

**Can you show me training data provenance documentation for your model?**

**Red flag answer:** Our training data is proprietary

**Good answer:** Shows documentation with dataset versions, sources, dates, composition

**Q: Walk me through how a genetic counselor would explain your AI's prediction to a patient**

**Red flag answer:** Talks about confidence scores and accuracy metrics

**Good answer:** Demonstrates actual explainability UI and reasoning pathway

**What happens when we need to retrain your model with updated ClinVar data?**

**Red flag answer:** "Models are pre-trained and validated, no retraining needed.

**Good answer:** Documents retraining protocol, timeline, validation requirements, costs

**If vendors can't answer these clearly with documentation and demos, they're selling research-grade AI, not production-ready clinical systems.**

---

Not sure what your results mean?

**Reach out to saurabh@nonstopio.com for a personalized AI readiness review.**