gart.

Self-Assessment: AI Infrastructure & Readiness

A practical diagnostic for CTOs, founders, and tech leaders

Before investing in AI tools, verify whether your infrastructure can support AI in production. This assessment helps identify gaps in data maturity, compute capacity, governance, and operational readiness — before AI initiatives stall or fail to scale.

⏱️ 5–7 minutes

📊 Instant results

🔍 Model-agnostic, implementation-focused

🎯 Built for teams deploying AI/ML in production

SECTION 1

Data Foundation & Quality

Q1. How accessible, structured, and reliable is the data needed to train or run AI models?

Choose one of the options, which is the most relevant to your case.

01

1) Fragmented & siloed data

- Data scattered across systems with no unified access
- Inconsistent formats, missing metadata
- Manual extraction required for most use cases
- High effort to prepare data for AI

02

2) Partially centralized, but inconsistent quality

- Some data lakes or warehouses exist
- Data quality issues are common (duplicates, gaps, drift)
- Limited documentation or lineage tracking
- AI teams spend most time on data prep

03

3) Mostly structured, with known quality gaps

- Core datasets are accessible and documented
- Data quality is monitored but not always enforced
- Some automation in place for ingestion and transformation
- AI-ready datasets exist for key use cases

04

4) Production-grade data infrastructure

- Unified data platform with clear access patterns
- Data quality, lineage, and governance are automated
- Real-time and batch pipelines are reliable
- AI/ML teams can self-serve high-quality datasets

Q2. How mature is your MLOps practice — from experimentation to production deployment?

Choose one of the options, which is the most relevant to your case.

1) Ad-hoc experimentation, no production path

- Models trained in notebooks, not versioned
- No clear path from experiment to deployment
- Production deployments are manual and fragile
- Model performance degrades unnoticed

2) Basic deployment pipeline, inconsistent practices

- Some models in production, but deployment is custom each time
- Limited monitoring or retraining workflows
- Model versioning exists but not enforced
- Rollback and testing are manual

3) Structured MLOps, with some automation

- Standardized deployment pipelines for most models
- Model versioning, monitoring, and logging in place
- Retraining workflows exist but may be manual
- Some drift detection and alerting

4) Fully operationalized ML lifecycle

- End-to-end automation from training to deployment
- Continuous monitoring, drift detection, and retraining
- A/B testing and rollback are standard practice
- Models are treated as production services

SECTION 2

Compute & Infrastructure Capacity

Q3. Do you have the compute resources needed to train, fine-tune, and serve AI models at scale?

Choose one of the options, which is the most relevant to your case.

1

1) No dedicated AI compute

- Models run on general-purpose infrastructure
- Training is slow or impossible for larger models
- No GPU or specialized hardware available
- Cost and performance are unpredictable

2

2) Limited or shared compute resources

- Some GPU access, but heavily contended
- Training jobs compete with other workloads
- Scaling is manual and expensive
- Inference latency is inconsistent

3

3) Dedicated AI infrastructure, but not optimized

- GPUs or TPUs available for AI workloads
- Some autoscaling and resource management
- Cost optimization is reactive
- Inference infrastructure exists but may bottleneck

4

4) Production-grade AI compute platform

- Right-sized compute for training and inference
- Autoscaling based on demand and cost targets
- Optimized for model serving (latency, throughput)
- Clear cost visibility and governance

Q4. How well do you monitor and understand AI model performance in production?

Choose one of the options, which is the most relevant to your case.

1) No production monitoring

- Models deployed but not tracked
- Performance degradation goes unnoticed

2) Basic logging, no actionable insights

- Logs exist but not analyzed
- No visibility into model drift or accuracy

3) Monitoring exists, but reactive

- Key metrics tracked (latency, accuracy)
- Alerts exist but investigation is manual

4) Proactive model observability

- Real-time drift detection and performance tracking
- Automated alerts with root cause context
- Model health dashboards for stakeholders

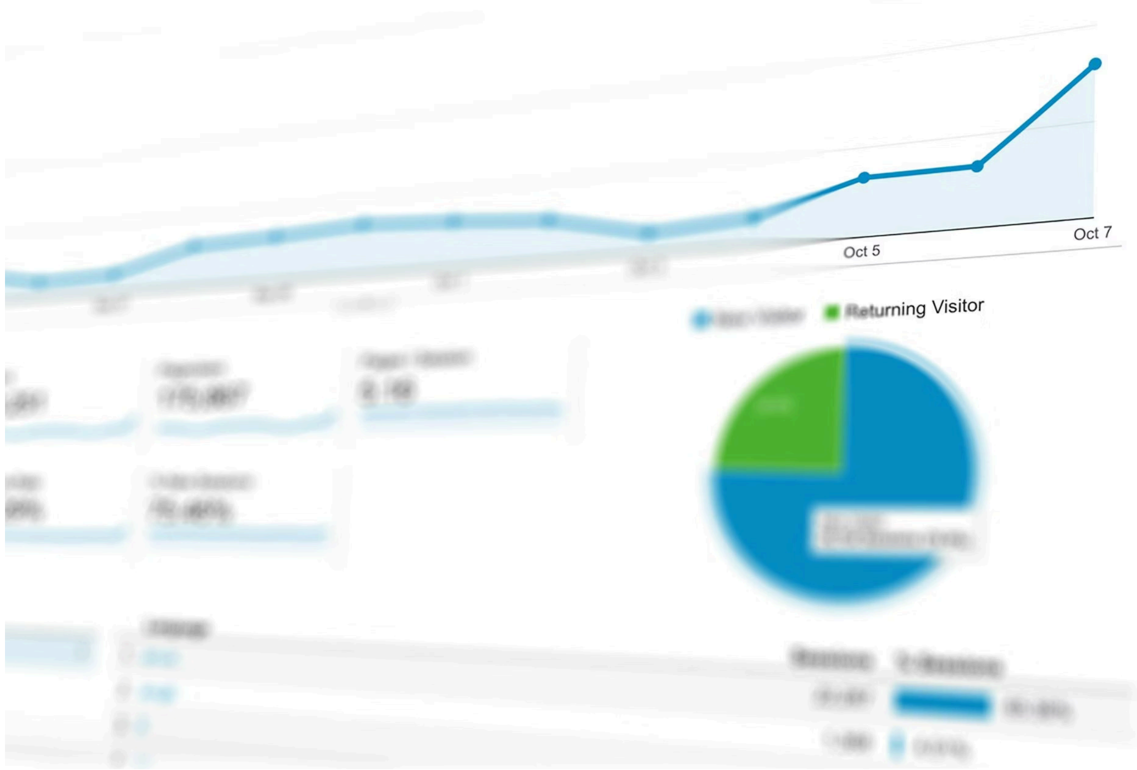

SECTION 3

AI Governance & Ethics

Q5. How are AI models governed for bias, fairness, explainability, and compliance?

Choose one of the options, which is the most relevant to your case.

1) No governance or oversight

- Models deployed without review
- High risk of bias or compliance violations

2) Informal guidelines, inconsistently applied

- Some awareness of risks, but no process
- Governance is reactive, not proactive

3) Governance framework exists, but not enforced

- Policies documented but not always followed
- Some bias testing and explainability tools in use
- Compliance checks are manual

4) Comprehensive AI governance in production

- Bias and fairness testing is automated
- Explainability and auditability are built-in
- Compliance requirements are tracked and enforced
- Clear accountability for model decisions

Q6. How well-equipped is your team to build, deploy, and maintain AI systems?

Choose one of the options, which is the most relevant to your case.

1) Limited AI/ML expertise

- Heavy reliance on external vendors or consultants
- No in-house model development capability

2) Some AI skills, but fragmented

- A few data scientists or ML engineers
- Knowledge is siloed, not shared across teams

3) Solid AI team, but gaps in production skills

- Strong experimentation and modeling skills
- Limited MLOps or infrastructure expertise
- Handoffs between teams are slow

4) Full-stack AI capability

- End-to-end ownership from research to production
- Cross-functional collaboration (data, ML, engineering)
- Continuous learning and upskilling culture

SECTION 4

Security & Privacy

Q7. How are AI models and training data protected from security and privacy risks?

Choose one of the options, which is the most relevant to your case.

1

1) No AI-specific security measures

- Training data includes sensitive information without controls
- Models are not secured or access-controlled

2

2) Basic security, but gaps remain

- Some data anonymization or access controls
- Model endpoints are protected but not hardened
- Privacy risks are not systematically assessed

3

3) Security and privacy controls in place

- Data access is controlled and audited
- Model endpoints have authentication and rate limiting
- Privacy impact assessments conducted for high-risk models

4

4) Comprehensive AI security posture

- End-to-end data encryption and access control
- Model versioning and provenance tracking
- Adversarial robustness and security testing
- Privacy-preserving techniques (differential privacy, federated learning) where needed

Q8. How well do you manage model versioning, reproducibility, and experiment tracking?

Choose one of the options, which is the most relevant to your case.

1) No versioning or tracking

- Experiments are lost or hard to reproduce
- No record of what was tried or why

2) Manual tracking in spreadsheets or docs

- Version history exists but incomplete
- Hard to compare experiments or reproduce results

3) Experiment tracking tools in use

- Tools like MLflow or Weights & Biases adopted
- Not all teams use them consistently
- Some gaps in reproducibility

4) Full experiment and model lifecycle management

- All experiments tracked with full context
- Models are versioned with lineage and metadata
- Reproducibility is guaranteed
- Easy rollback and comparison across versions

SECTION 5

Cost & ROI

Q9. How well do you understand and manage the costs of AI infrastructure and operations?

Choose one of the options, which is the most relevant to your case.

1	**No visibility into AI costs** • Training and inference costs are unknown until bills arrive • No budget allocation or accountability • Runaway spending is common
2	**Basic cost tracking, no optimization** • Can see total AI spend but not per-model or per-team • No cost forecasting or budgeting • Optimization is reactive
3	**Cost visibility with some controls** • Costs tracked by model, team, or project • Budgets and alerts exist • Some optimization (instance types, batch sizes) • ROI is estimated but not systematically measured
4	**Full cost governance and ROI tracking** • Real-time cost visibility per model and workload • Automated cost optimization (spot instances, autoscaling) • Clear ROI metrics tied to business outcomes • Continuous cost-performance optimization

Q10. How well does your AI infrastructure scale with model complexity, data volume, and user demand?

Choose one of the options, which is the most relevant to your case.

1

Scaling is manual and unpredictable

- Training larger models requires significant re-architecture
- Inference can't handle traffic spikes
- Performance degrades under load

2

Scaling works but requires heavy engineering

- Can scale with effort and planning
- Frequent bottlenecks and performance issues
- Cost increases faster than value

3

Scaling is mostly automated but not optimized

- Autoscaling exists for training and inference
- Some inefficiencies remain (over-provisioning, latency)
- Can handle growth but not always cost-effectively

4

Elastic, efficient, and production-ready scaling

- Seamless scaling from prototype to production
- Optimized for cost and performance
- Handles traffic spikes and model updates gracefully
- Supports rapid experimentation and iteration

How to Use This Assessment

For each question, select the option that best describes your current state. Be honest — this is a diagnostic tool, not a test. Each option is worth points:

Level 1: 1 point	**Level 2: 2 points**	**Level 3: 3 points**	**Level 4: 4 points**

Add up your total score across all 10 questions to determine your AI readiness level. The interpretation guide will help you understand what your score means and what to do next.

Self-Assessment Result Interpretation

🔴 0–10

AI Experimentation Stage

What this means:

- AI is exploratory, not production-ready
- High risk of failed pilots or stalled initiatives
- Infrastructure gaps will block scale

Recommended next step: AI Readiness Assessment & Infrastructure Planning

🟠 11–20

Early AI Adoption

What this means:

- Some AI models in production, but fragile
- Data, governance, or MLOps gaps create risk
- Scaling will be difficult without investment

Recommended next step: MLOps Foundation & Data Platform Build

🟢 21–30

Production AI Capability

What this means:

- Solid foundation for AI in production
- Some optimization and governance gaps remain
- Ready to scale with targeted improvements

Recommended next step: AI Governance, Cost Optimization & Advanced MLOps

🔵 31–40

AI-Native Organization

What this means:

- End-to-end AI capability from research to production
- Infrastructure supports rapid iteration and scale
- Ready for advanced AI use cases and competitive advantage

Recommended next step: Continuous Innovation & AI Excellence Programs gart.

At Gart Solutions, we help teams move from AI experimentation to production-ready AI infrastructure.

Ready to build a concrete AI roadmap?

Start an AI Readiness Audit

Book an AI Infrastructure Assessment Call