



Assess your storage strategy for the AI era

Most storage strategies weren't built for the scale, speed, and demands of the AI era, and the gap is starting to show.


Assess your readiness to scale AI and build a data-centric storage strategy that turns infrastructure into strategic advantage.

Time to complete: 5-7 minutes

Storage strategy is the foundation of AI value creation

Enterprises are generating and consuming more data than ever before, spanning everything from operational telemetry and transaction records to customer interactions and machine-generated insights.

However, many are unable to fully capitalize on the value of their data because legacy storage strategies were not designed to scale for the demands on AI and are rapidly becoming outdated.



This tool helps you assess whether your current storage strategy actively enables AI-driven value creation or quietly constrains it.

Shift to a data-centric operating model

As AI moves from pilot to production scale, storage strategies must evolve—aligning architecture, operations, and investment decisions around data as the focal point of value creation.

Strategy starts with data gravity

AI is changing how data flows, distributed across edge, core, and cloud. Storage strategies must align to where data is actively created and used, not where it historically lived.

Data availability shapes AI velocity

AI depends on consistent, reliable access to data. Modern storage strategies enable dynamic movement and tiering, balancing throughput, capacity, and governance requirements.

Storage density is the currency of AI

As AI adoption accelerates, organizations become limited by power and space constraints—making efficiency metrics like terabytes per watt and per square foot central to long-term planning.

Strategic planning for what's next

AI demand is not static. Leading strategies look ahead—aligning storage investments to AI roadmaps, anticipating future demand, and building flexibility into long-term capacity planning.

Assess your current storage state



Evaluate your organization across four dimensions. For each question, select the answer that best reflects your current state.

Section 1: Data gravity

1 | How do you account for where data is created and consumed by AI systems?

- Most data is moved to a primary cloud or compute-centric neocloud
1 point
- Data location is considered for certain data types, but AI/ML activities are typically centralized
2 points
- Storage and compute are aligned to where data is generated, its value, and access times across environments
3 points

2 | How is data gravity factored into infrastructure decisions?

- Infrastructure decisions are driven by standardization, cost, or historical patterns
1 point
- It's acknowledged in certain projects but isn't consistently embedded in architectural decisions
2 points
- It's a core design principle, shaping how we architect storage, compute, and data pipelines
3 points

3 | How well does your strategy scale with increasing data gravity?

- Our strategy becomes harder to scale as data volumes grow
1 point
- We can scale in controlled environments, but growing edge and distributed data sources introduce friction
2 points
- Our strategy is designed to scale with data gravity, supporting distributed, high-volume data without compromising agility or outcomes
3 points

DATA GRAVITY SUBTOTAL /9

Section 2: Data availability

4 | How do you prioritize data availability for AI and intelligent decision-making?

- Availability is optimized primarily for retention cost
1 point
- Cost and availability are balanced for priority workloads, but access is not consistently optimized across datasets
2 points
- Time-to-data is a critical consideration enabling rapid access to both real-time and long-term context
3 points

5 | How has your storage tiering approach evolved for AI workloads?

- Traditional hot and cold tiers, optimized for retention cost
1 point
- Evolved beyond basic tiering, with selective optimization for active datasets
2 points
- Agent-aware tiering that prioritizes data retrieval, continuity, and preserving long-term context
3 points

6 | Can you place data on the most efficient storage tier without impacting AI performance?

- Architectural limitations force us to keep most AI data on expensive, high-performance tiers
1 point
- Some data is placed on capacity-optimized tiers, but data movement and access constraints limit their use for active workloads
2 points
- Our architecture enables high-throughput, parallel data movement across tiers without compromising performance or scalability
3 points

DATA AVAILABILITY SUBTOTAL /9

Section 3: Infrastructure mix & density

7 | How do you scale storage capacity?

- By adding racks, hardware, or facility footprint
1 point
- By combining footprint expansion with incremental density improvements
2 points
- Through density-led design that increases capacity within existing power and space budgets
3 points

8 | How do you optimize storage ROI today?

- Basic metrics like storage capacity and utilization
1 point
- Capacity metrics, plus cost and performance per terabyte across workloads and environments
2 points
- Advanced metrics including cost and power efficiency and AI/ML-driven business performance improvements
3 points

9 | How do you view storage density?

- Primarily as an engineering or hardware optimization
1 point
- Important for efficiency but not a major planning factor
2 points
- As a strategic planning lever that directly shapes AI value creation and long-term economic upside
3 points

INFRASTRUCTURE MIX & DENSITY SUBTOTAL /9

Section 4: Capacity planning horizon

10 | How far ahead do you plan storage capacity?

- Quarterly or annual budget-driven planning
1 point
- 12- to 24-month forecasts loosely tied to business initiatives
2 points
- 3+ year planning aligned to AI, data growth, and business strategy
3 points

11 | What typically triggers storage procurement?

- Systems approaching capacity limits or unexpected shortages
1 point
- Forecasted growth influenced by near-term utilization
2 points
- Demand is modeled for high-value data sets, data expansion, and AI workloads
3 points

12 | How well do you forecast AI storage needs?

- AI workloads are unpredictable and not factored into capacity planning
1 point
- Some AI workloads are included but forecasting is inconsistent or short-term
2 points
- AI workload growth is modeled into multi-year capacity planning
3 points


CAPACITY PLANNING SUBTOTAL /9

Calculate your storage strategy state



Add your scores to see where your enterprise falls on the storage strategy spectrum.

TOTAL
SCORE



12-20 points: Reactive operator

Storage is operating as a constraint—not a strategy.

What this means:

Data storage at your organization is likely treated as an operational overhead cost focused on immediate needs. Reactive decisions around capacity and data retention make it harder to keep up with data growth and AI demands. What looks manageable today can quickly become a bottleneck tomorrow.

What to do next:

Shift from reactive to forward-looking. Build a multi-year capacity plan, factor in power and density constraints early, and align data retention with the long-term AI needs of your enterprise. Treat storage not just as infrastructure, but as a driver of future data value.

- Establish a baseline: quantify current capacity, constraints, and retrieval performance.
- Identify AI-critical datasets and ensure their timely accessibility.
- Understand AI/ML data growth needs and related cost-benefits.
- Create alignment between infrastructure, AI, and finance stakeholders.

21-28 points: Transitional planner

Storage is evolving—but not yet enabling full-scale AI performance.

What this means:

You've started to think strategically about storage, but the elements are not fully connected yet. Pockets of long-term planning exist, but decisions around procurement, scaling, and data retention are still sometimes reactive. AI, infrastructure, and finance aren't always moving in sync.

What to do next:

Bring the pieces together. Build a 3- to 5-year roadmap that links AI/ML objectives and related workloads to storage capacity, density, cost, and supply planning. This will ensure your infrastructure keeps pace with AI/ML-based productivity increases.

- Consolidate AI/ML, infrastructure, and business forecasts into a single planning model.
- Define clear service-level agreements (SLAs) for data retrieval and availability.
- Identify where density improvements can reduce dependency on physical expansion.

29-36 Points: Strategic infrastructure leader

Storage is a strategic asset—and a source of competitive advantage.

What this means:

Storage is a core part of your long-term business strategy. Decisions around data gravity, capacity, density, and data availability are tightly aligned with AI goals and capital planning. At this stage, the focus shifts to sustaining and extending advantage.

What to do next:

Keep your edge. Stress-test your current roadmap against accelerating AI growth, tightening power constraints, and ongoing market volatility to ensure your infrastructure continues to scale without compromise.

- Benchmark density, power efficiency, and cost performance.
- Identify emerging constraints in data accessibility and supply.
- Fine tune AI/ML use cases and align infrastructure evolution to next-gen AI architectures and workloads.

Risk & exposure mapping



Use the checklists below to identify emerging risks and exposures, helping you prioritize where to take action against potential roadblocks.

Where are hidden constraints forming?

In many organizations, multiple issues exist simultaneously, often reinforcing one another. Mapping these risks will help you start taking targeted action that will have the greatest impact.

Capacity risks

If data storage planning is short-term or storage expansion is reactive, you may be experiencing:

- AI rollout delays when storage can't keep pace with deployment
- Cost spikes from emergency storage procurement
- Data growth exceeding planned storage capacity

Bottom line: Storage becomes a bottleneck that slows data-driven initiatives and AI deployment and growth.

Power & density risks

If storage growth is primarily achieved through adding hardware and racks rather than improving density and power efficiency, you may be encountering:

- Data center power limits that constrain infrastructure expansion
- Rising operating costs due to inefficient power usage
- Data center footprint expansion that increases facility and infrastructure costs

Bottom line: Without gains in density and power efficiency, scaling storage becomes harder and more expensive over time.

AI execution risks

If archival data is difficult to retrieve, or storage architecture is not aligned with AI workflows, you may be experiencing:

- Limited access to historical data for AI model retraining and experimentation
- Slow data retrieval that hinders AI model development
- Past data deletions restricting AI capabilities and future potential

Bottom line: Your enterprise data exists, but on tape or cold cloud storage, it probably isn't usable at the speed or scale that AI requires.

Ready to rethink your storage strategy?

The ability to scale intelligence depends on your ability to scale storage. With the right partner, storage becomes a source of competitive advantage.

No matter where you are in your storage maturity journey, Seagate brings the technology, expertise, and scale to help you move from managing storage to maximizing its impact.

[Talk to an expert to start building your roadmap](#)



SEAGATE