

More capacity. Less infrastructure.

Executive summary

Accelerating data growth is pushing storage deployments to their practical limits. Rack space, available power, and cooling capacity increasingly determine what can be deployed efficiently, shifting storage planning from individual component selection to system-level behavior within fixed physical and electrical constraints.

This brief evaluates those system-level effects directly. Rather than comparing drives in isolation, it examines rack-level storage architectures. It answers a practical planning question: how much usable capacity can be delivered within a fixed amount of rack space, and how do power and energy requirements scale with capacity?

The analysis shows how differences in solid state drives (SSDs) and hard disk drives (HDDs) translate into rack count, floor-space requirements, and aggregate drive energy savings, using a realistic assumption of 36U of usable rack space (which reserves space for networking, power distribution, and cooling in a standard 42U rack). The 245.76TB Micron 6600 ION SSD demonstrates this advantage by delivering substantially higher capacity per rack than conventional 36TB HDD-based designs, enabling the same capacity targets to be met with less infrastructure and lower energy consumption. These differences appear at all deployment sizes and become increasingly material as deployment capacity grows.

The comparison that follows is based on a consistent rack-level modeling framework using publicly available per-drive measurements. Table 1 summarizes the resulting first-order space and energy effects at a representative 100PB deployment, establishing a baseline for the scaling analysis that follows.

Metric	Micron 6600 ION-based storage ¹	36TB HDD-based storage ²
Deployed capacity	100PB	100PB
Racks required	1 rack	≈ 4 racks
Capacity per rack	≈ 176.9PB	≈ 25.9PB
Relative data center footprint	1x	4x
Annual drive energy for an example 70% activity ³	≈ 75,000 kWh/year	≈ 170,000 kWh/year
Scaling behavior	Capacity grows within existing racks, preserving space and power headroom	Capacity grows by adding racks, servers, and supporting infrastructure

Table 1: 100PB example racks, footprint, energy use, and scaling behavior summary

1. Micron 6600 ION E3.L 245.76TB SSD capacity per rack calculated as: 2U server, 40x 245.76TB SSDs per server, 18 servers per rack; 40 * 245.76TB per SSD * 18 servers ≈ 176.9PB per rack. SSD energy calculated assuming 100PB deployed capacity using 407x 245.76TB Micron 6600 ION SSDs drawing 70% maximum rated power (of 30 watts each) in use 24 hours a day for one year = (407 * 0.7 * 30 * 8760) ≈ 75,000 kWh/year. DOE/NREL guidance indicates that, while total rack area is higher, net IT white space planning typically allocates ≈ 20 sq ft per rack, with the remaining space consumed by aisles, power, and cooling. See "Best Practices Guide for Energy-Efficient Data Center Design" (energy.gov) for additional information.

2. HDD capacity per rack calculated as: 5U server, 100 36TB HDDs per server = 3,600TB per 5U server = 720TB per U. Capacity per rack = 720TB per U * 36U = 25.9PB/rack. Note: While a 36U rack can house only 7x 5U servers, this extension is used solely to normalize capacity per rack for first-order comparison; it intentionally favors the HDD case to avoid overstating SSD density advantages. HDD energy calculated assuming 100PB deployed capacity using 2,778x 36TB data center HDDs drawing 70% of their maximum rated power (of 10 watts each) in use 24 hours a day for one year = (2,778 * 0.7 * 10 * 8760) ≈ 170,000 kWh/year. For HDD-based configurations, ≈ 20 sq ft per rack is also assumed (from the same source). While future HDD capacity increases may improve absolute HDD density, they do not change the order-of-magnitude rack-level gap under current form-factor constraints. Activity level means the percentage of maximum drive power. A 70% activity level is an example only; different values may produce different results, as shown later in this document.

Introduction

To understand the real impact of storage density on data center planning, the analysis below demonstrates how different architectures scale in space and power. Rather than focusing on individual drives, it evaluates complete rack-level configurations and the resulting footprint and energy requirements as capacity grows.

We focus on the metrics that data center leaders demand when planning storage infrastructure: the number of racks required (for fixed capacities), capacity per rack, watts per provisioned PB, and annual kWh for representative deployments.

Micron 6600 ION SSD configuration		Data center HDD configuration	
Rack space	36U	Rack space	36U
Chassis size	2U	Chassis size	5U
Drives per chassis ⁴	40 (2U chassis)	Drives per chassis	100 (5U chassis)
Capacity per drive	245.76TB	Capacity per drive	36TB
Capacity per U ⁵	4.92PB	Capacity per U	0.72PB

Table 2: SSD and HDD configurations

With these configurations established, the analysis now examines how each architecture translates rack space into usable capacity as deployments scale.⁶

Figure 1 demonstrates the magnitude of the storage density advantage when moving from HDD-based designs (shown in grey) to Micron 6600 ION SSD configurations (shown in purple). As rack space is incrementally filled, the Micron 6600 ION SSD configuration delivers more capacity per occupied unit than the HDD configuration at every step, resulting in higher usable capacity per rack.

Furthermore, by delivering far higher capacity per rack and per chassis, SSD-dense systems reduce footprint and management overhead while simplifying architectural right-sizing, because their finer step sizes enable capacity to be added in smaller, more precise increments that reduce over-provisioning, minimize stranded resources, and improve long-term planning economics.

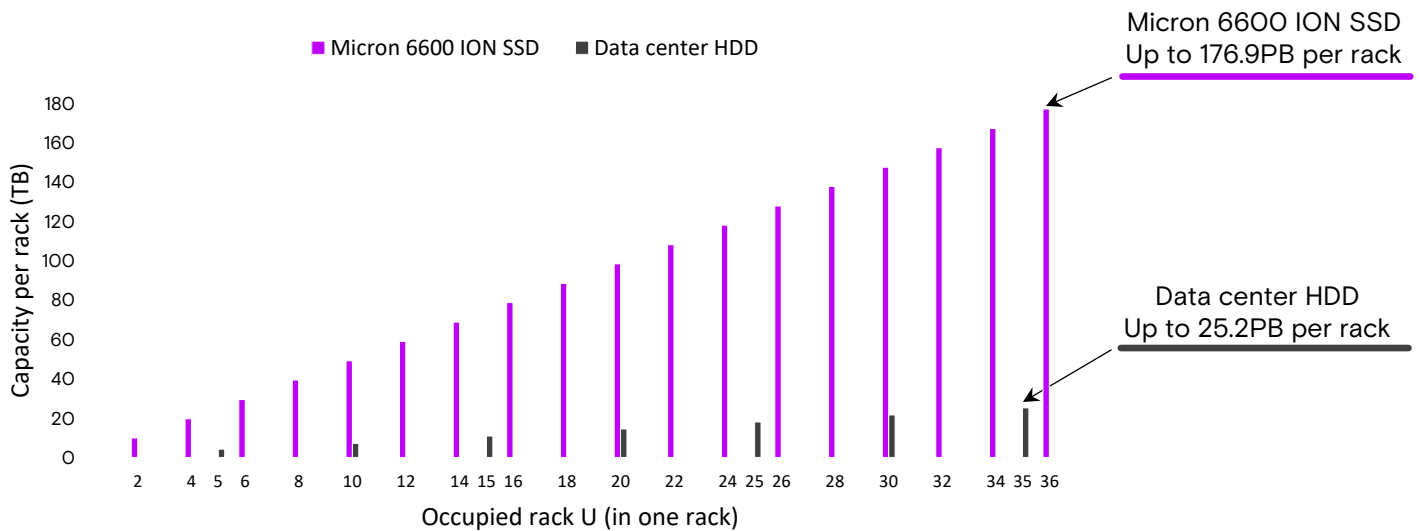


Figure 1: Provisioned capacity versus occupied rack space per rack (higher is better)

3. Drives per chassis values are based on public information available at the time of this document's initial publication.
 4. Micron 6600 ION SSD capacity per chassis calculated as: 40x SSDs per chassis, chassis capacity = 245.76TB * 40 SSDs = 9.83PB in 2U = (9.83 / 2) ≈ 4.92PB per U. HDD capacity per chassis calculated as: 100 HDDs per 5U chassis = 36TB * 100 HDDs = 3.6PB per 5U chassis = 0.72PB per U.
 5. The number of racks needed is calculated by dividing the provisioned capacity by the capacity per rack and rounding up to the next whole rack count.

Rack count implications of capacity density

As rack-level capacity increases, the most direct planning consequence is a reduction in the number of racks required to deliver a given amount of storage. Because higher-density configurations convert each occupied rack unit into more usable capacity, total rack count grows more slowly as provisioned capacity increases.

This relationship between capacity density and rack count is independent of assumptions about absolute rack height. Under typical 36U deployments, the relative differences between SSD and HDD designs remain consistent. Higher storage density means fewer racks are needed to reach the same capacity target, reducing overall footprint, infrastructure complexity, and associated capital requirements.

Figure 2 illustrates how the rack-level capacity behavior shown in Figure 1 translates into the number of racks required to meet specific provisioned capacity targets. The SSD (in purple) and HDD (in grey) indicators show how differences in capacity density directly affect rack counts as deployments scale.

The rack-count growth behavior shown in Figure 2 illustrates how differences in capacity density directly drive rack requirements and floor-space needs. As storage capacity increases, SSD-based configurations require significantly fewer racks to reach the same targets, reducing floor-space requirements and the associated infrastructure that scales with rack count.




Figure 2: SSD and HDD rack requirements as provisioned capacity scales (lower is better)

Table 3 presents the same rack-count data shown in Figure 2, expressed as a comparison across representative deployment sizes. This tabular view complements the chart by making the absolute rack requirements explicit for common capacity targets, supporting straightforward footprints and planning comparisons.⁷







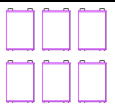
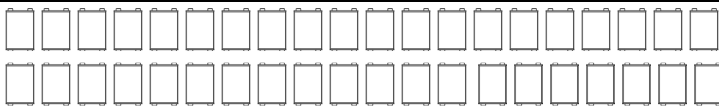
Provisioned capacity	SSD		HDD	
	Racks	ft ²	Racks	ft ²
200PB		40		160
300PB		40		240
500PB		60		400
1000PB		120		800

Table 3: Rack count required by provisioned capacity. Each rectangle represents one rack.

6. Table 3 represents the data in Figure 2 as the necessary floor space for four example provisioned capacity points assuming 20ft² per rack as noted earlier. Other provisioned capacity points and/or floor space allocations will yield different results.
 7. Power scaling reflects instantaneous drive power only. All power scaling calculations assume that 1x SSD consumes 30 watts and each HDD consumes 10 watts.

Power scaling behavior by provisioned capacity

Figure 3 shows the instantaneous drive power required (kW) as provisioned capacity increases from an example 10PB to 1EB of provisioned capacity (additional savings may be realized through infrastructure reductions). Across these capacity points, the SSD-based configuration consistently demonstrates lower power requirements per provisioned petabyte than the HDD-based configuration, with SSD power (purple) remaining below HDD power (grey) at the shown provisioned capacity.⁸

This behavior reflects how higher storage density can affect power scaling at the system level.

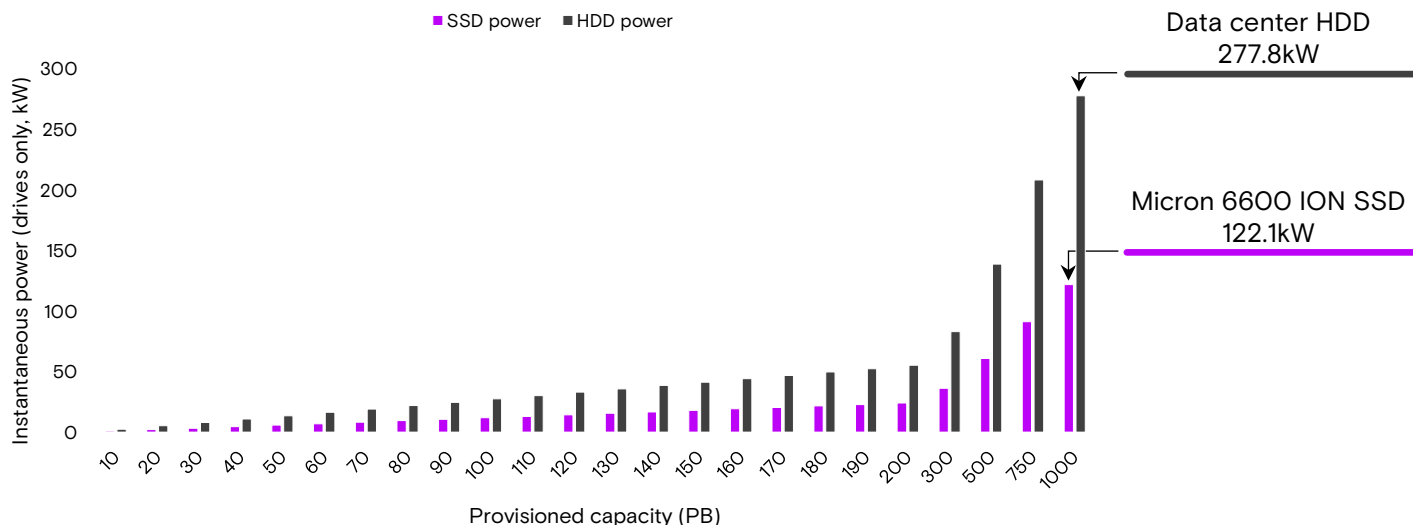


Figure 3: SSD and HDD instantaneous power needed by provisioned capacity (lower is better)

Annual energy scaling by provisioned capacity and activity level

For large-scale storage deployments, yearly energy consumption is a stable, comparable measure of storage efficiency. While instantaneous power use gives a snapshot, annualized energy captures its cumulative effect on power budgets, operating costs, and sustainability metrics.

To reflect realistic deployment scenarios, this analysis evaluates four provisioned capacities—200PB, 300PB, 500PB, and 1EB—spanning enterprise, cloud, and hyperscale environments. At 200PB, storage density and power can influence architectural decisions. At 300PB and 500PB, multi-rack efficiency and floor-space considerations become increasingly important. At 1EB (1000PB), the combined effects of density, rack count, and power scaling help provide a clear view of the large-scale infrastructure implications.

Because storage systems rarely operate at a single, fixed activity level, the analysis models annual energy across activity levels ranging from 10% to 100% of maximum storage device I/O utilization. This approach captures both steady-state and peak-demand conditions, ensuring the results reflect day-to-day operation as well as periods of elevated load that influence system sizing, cooling design, and long-term operating cost.

Figure 4 shows yearly drive energy consumption (kWh) as a function of storage activity level for four provisioned capacities ranging from 200PB to 1EB. For every capacity and activity level shown, the SSD-based configuration consumes less annual energy than the HDD-based configuration, with SSD curves consistently below their HDD counterparts.

This annual view captures the cumulative effect of the power-scaling behavior shown in Figure 3. As higher-density SSD configurations require less power per provisioned petabyte, those differences compound over time into materially lower yearly energy consumption as deployments scale.

8. Power scaling reflects instantaneous drive power only. All power scaling calculations assume that 1x SSD consumes 30 watts and each HDD consumes 10 watts.

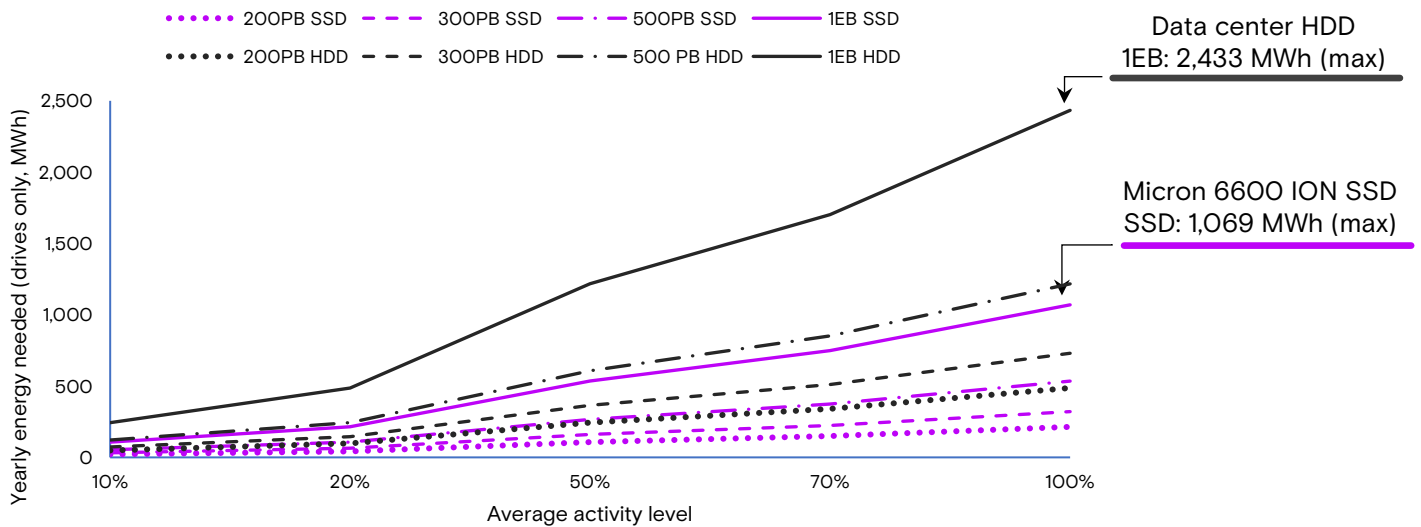


Figure 4: SSD and HDD yearly power needed by provisioned capacity (for each provisioned capacity, lower is better)

For planners and operators, the implications are measurably reduced operating energy use, lower cooling demand, and improved sustainability outcomes—benefits that become more pronounced as both capacity and utilization increase.

Figure 4 shows that, at every modeled capacity and activity level, the Micron 6600 ION SSD configuration consumes less annual drive energy than the HDD configuration, maintaining a lower energy profile even as total energy use increases with scale.

Infrastructure planning implications at scale

Translating technical gains into business outcomes helps organizations connect infrastructure efficiency to financial and operational strategy. Converting yearly drive-energy estimates into costs using publicly available U.S. or global energy-rate data (such as [this document from eia.gov](#)) provides a practical way to relate energy consumption to energy-related operating expenses and to express capacity decisions in financial terms.

For business leaders, this standards-based conversion provides a consistent, standards-based financial framework for evaluating rack-space deferral, infrastructure simplification, and sustainability considerations.

For data center planning engineers, it provides a consistent methodology for translating per-drive measurements into rack-level power and energy models.

Together, these elements form a coherent framework that links deployment-level architecture choices to measurable business outcomes while preserving engineering transparency and rigor.

- Floor-space deferral: Fewer racks can defer or eliminate costly expansion projects; a simple proxy of ≈ 20 sq ft per rack highlights potential savings at scale.
- Lower installed-base complexity: Each avoided rack often removes at least one top-of-rack switch, PDUs, and structured cabling, reducing both upfront capital outlay and long-term operational overhead.
- Sustainability: Higher density reduces infrastructure overhead per TB, while SSD performance reduces the time systems operate in higher power states, lowering both energy consumption and cooling load for equivalent workloads.

Conclusion

As storage capacity continues to grow, data center planning is increasingly shaped by practical limits rather than theoretical ones. Rack space, available power, and operational complexity all influence what can be deployed efficiently. This analysis evaluates those factors using a realistic 36U of rack space and compares two capacity-optimized storage architectures under consistent assumptions.

Within that framework, Micron 6600 ION SSD systems provide materially higher usable capacity per rack than 36TB HDD-based designs. That difference reduces the number of racks required to support provisioning targets across the evaluated range. Fewer racks translate into a smaller footprint, with effects that become more pronounced as capacity scales.

Power analysis shows a similar trend. Based on publicly available per-drive estimates and extended to rack-level systems, the SSD configuration demonstrates lower power requirements across the modeled capacities. When expressed as annual energy consumption, these differences help reduce operating energy, an increasingly relevant consideration in power-limited environments.

SSD-dense architectures deliver measurable planning advantages. Higher per-chassis capacity and smaller scaling increments enable finer-grained capacity addition, reducing the likelihood of forced overprovisioning and stranded resources. This can simplify deployment planning while maintaining flexibility for future growth within existing facilities.

Overall, the results illustrate how storage density influences rack count, floor space, and power demands under realistic deployment conditions. When evaluated at the system and rack level, Micron 6600 ION SSDs support denser configurations with fewer racks and more predictable scaling characteristics, providing a practical option for organizations planning for continued capacity growth.

Organizations evaluating large-scale storage architectures can find additional technical specifications, configuration options, and supporting resources at micron.com/6600.

Methodology and sensitivities

Capacity calculations are based on decimal TB/PB to align with public datasheets (36TB, 245.76TB, 100PB, etc.). In contrast, the total available storage space per rack is set at 36U for this analysis (while data center racks are commonly 42U). This rack space difference is designed to reserve space for top-of-rack networking, bottom-of-rack power, etc., to better align with a production deployment.

Per-drive power consumption is derived from publicly available, open-access, per-drive measurements for capacity-focused data center SSDs and HDDs, like [this Micron product brief](#) and [this Seagate product manual](#). Facility energy is estimated by multiplying IT energy by PUE. This analysis uses a representative PUE of 1.4. Organizations should substitute their facility's actual PUE for site-specific results.

How we model the real-world footprint and energy

Footprint model: ≈ 20 square feet per rack (aisles + service clearance) to translate rack count into data center floor space.

Energy model: Publicly reported per-drive measurements for high-capacity data center SSDs and HDDs, normalized to system and per-PB values for apples-to-apples comparison.

micron.com/6600-ION