

Micron® Flex Capacity™ Feature — Tunable Storage in a Flash

Micron® 9200 Series of NVMe SSDs and Micron® 5100 SATA SSDs
Give You Control of Performance and Capacity

Fast Facts

- Unleash Your Applications: Easily match SSD performance and capacity to your needs.
- Lower Complexity: Manage fewer drive types and enable better overall economics.
- Reduce Risk: Qualify one NVMe or SATA SSD, tune it for multiple workloads.
- Reduce Fees: Control the Micron Flex Capacity feature through Micron Storage Executive software with no license fees.



Meet the needs of a variety of applications, workloads, and requirements with the Micron Flex Capacity feature.

Overview

The advanced Micron® FlexPro™ SSD firmware architecture is designed to unleash the true capabilities of the storage media, giving IT administrators the ability to tune the drive so that it meets specific workload characteristics such as performance, endurance and security. First introduced in the Micron® 5100 SATA portfolio, followed quickly by the Micron® 9200 Series of NVM Express® (NVMe™) SSDs, the Micron FlexPro architecture provides a consistent experience across all Micron SSDs. The Micron FlexPro architecture creates commonality and gives longevity to SSD investments, dramatically simplifying the selection, qualification and deployment processes.

The Micron® Flex Capacity™ feature (part of the Micron FlexPro architecture) enables easy adjustments to the Micron 5100 SATA SSDs and Micron 9200 SSDs with NVMe product capacity to match a wide variety of applications, workloads, and requirements — all serviced from fewer drive stock keeping units (SKUs) for better overall economics and lower support costs.

As more storage architects are turning to SSDs, there is a commensurate demand for “tuned” storage capacity and fills per day for several reasons. Some deployments demand an SSD with greater capacity (where the workload is highly read-focused), some demand a cost-sensitive approach to workloads that need read/write balance, and some need an SSD that offers greater write performance and fills per day. Budgets demand that we satisfy both capacity and fills-per-day needs across a broad range of interfaces and that we do it with the smallest number of different SSDs possible to minimize qualification and inventory expense.

Balancing these competing demands has always been a challenge. We have to choose different types of SSDs for our read-focused workloads, our write-focused workloads and our balanced read/write use needs — and we have to qualify each type separately, as well as stock appropriate spares for long-term maintenance. Although this tactic of buying specialty SSDs for each workload can produce better results for each application, the number of different SSDs to buy, stock, deploy, and manage can be taxing — and each may require a compromise (for example, capacity or fills per day is optimal).

This tech brief describes how the Micron 5100 enterprise SATA SSD's Flex Capacity feature enables precise, application-level optimization while also overcoming the challenges of performing multiple SSD qualifications and storing excessive inventory spares (SKUs) for maintenance.

The Micron Flex Capacity feature enables easy adjustments to the Micron 5100 SATA SSDs' and Micron 9200 Series of NVMe SSDs' capacity and write performance while enabling easy tuning to match a wide variety of applications, workloads, and requirements — all serviced by fewer drive SKUs for better overall economics and lower support costs.

What Is Over-Provisioning?

Over-provisioning (OP) on SSDs refers to spare capacity that the drive may use to optimize internal processes. All flash-based SSDs contain some level of OP. For example, suppose we have an SSD with 1TB of NAND on the circuit board (1TB equals 1024GB), which is the "raw" capacity of this SSD. Typically, having a terabyte of storage on the circuit board does not mean that the entire terabyte is available for system use.

Depending on the SSD design, 960GB, 800GB, or 480GB of the 1TB of space may be available. Each of these SSDs would have different advertised capacities — 960GB, 800GB, or 480GB, respectively. We would typically think of each as a different product.

Figure 1 shows how OP is calculated for SSDs. The blue area is the total media capacity on the SSD circuit board (1TB in this example). The green

crosshatched area is the portion that is available for the system to use. The difference is the drive's OP, which depends on the drive's design.




Figure 1: OP of an SSD

Over-Provisioning, Write IOPS Performance, and Capacity

SSDs are designed with some level of OP, which affects the SSD's write IOPS performance and available capacity. Using our 1TB raw capacity example mentioned in the previous section, we can set the OP level very low and make 960GB of the 1TB available. An SSD like this is well suited for read-intensive storage (static data that rarely changes but is often read, like in catalogs, data archives, capacity/bulk storage, etc.). This 960GB SSD would have great read performance but limited write performance due to its lower OP. The OP level is usually set in the factory. Modifying the factory OP level used to be difficult.

Changing the OP level used to be so difficult that if we needed an SSD designed for small I/O size mixed-use (like 8K 67% read, 33% write), the manufacturer could take that same 1TB raw capacity SSD and make more OP available, setting the user capacity to 800GB, for example. The extra OP would give this 800GB drive better write performance — but less user capacity. This too would be set at the factory and would not be easy to modify.

Lastly, if we needed a write-intensive SSD, the manufacturer might take the same 1TB raw design and set the OP to make 480GB available. This SSD (with still more OP) would have the greatest write performance of all three examples, but the smallest user capacity.

SSD designers have controlled these compromises for many years. For a given raw SSD, the more OP the factory sets, the more write IOPS performance but the less available capacity, which is a tradeoff.

Micron Flex Capacity Feature Puts *You* in Control

With the Micron Flex Capacity feature, you can easily optimize a Micron 5100 or Micron 9200 drive for diverse applications and workloads, managing the exact write performance and capacity you need — enabling you to source fewer types of SSDs while matching the SSD’s capabilities to your needs.

Figure 2 shows how we can use the Micron Flex Capacity feature to improve the SSD’s write performance while simultaneously adjusting its available capacity.




Figure 2: Flex Capacity Feature at Work

Note that the raw media capacity remains the same for each drive — only the available capacity (green) changes as we apply the Micron Flex Capacity feature to tune the SSDs.

Figure 2 shows the effect of the Micron Flex Capacity feature on available space. On the left is an example Micron 5100 SSD that is factory-configured with an available 960GB of capacity. At 960GB, this Micron 5100 SSD offers the Micron default user/system capacity, the factory-set over-provisioning and the factory-set write IOPS performance. The factory-set capacity of 960GB is a good default option; it works

well in a variety of enterprise applications, workloads, and deployments.

The center drive in Figure 2 shows how we can use the Micron Flex Capacity feature to improve the Micron 5100 SSD write performance and slightly reduce the available capacity to 800GB. The advantages are twofold: The 800GB user capacity enables better write IOPS performance, or if we need to replace an 800GB SSD (perhaps from another vendor or because one has failed), we can use the 960GB Micron 5100 SSD and easily reset it to 800GB for an exact match.

The drive shown on the right in Figure 2 shows another option where we can use the Micron Flex Capacity feature to further improve write performance by, in turn, tuning the user capacity to 480GB. We can now use the same Micron 5100 SSD for even more write-intensive workloads or use it to replace a 480GB SSD that may have failed.

In this example, we reset a 960GB Micron 5100 SSD to tune write IOPS performance or match the capacity of an existing 800GB or 480GB SSD. However, with the Micron Flex Capacity feature, it is just as easy to reset the Micron 5100 SSD to 627GB, 472GB, or any capacity needed less than the factory-set (default) capacity.

The same effect is available for the Micron 9200 SSD, albeit with different factory and user capacities.

The Micron Flex Capacity feature can also help minimize the impact on qualification resources. One SSD SKU of a product family can be qualified for read-intensive workloads. The same SKU can be set to a slightly reduced capacity for mixed use or set to a further reduced capacity if the workload is write-intensive. The same SSD and same qualification has multiple uses.

Use the Micron Flex Capacity Feature to Tune Write IOPS Performance

Because the Micron Flex Capacity feature lets you easily change the available gigabytes, you can also choose to make write IOPS performance and capacity adjustments permanent or temporary. You

can set your SSD to best match known workloads or change the characteristics to manage unexpected application I/O demand more easily. You can increase the SSD's write IOPS performance permanently or only when you need it, for as long as you need it. The Micron Flex Capacity feature lets you choose what is best.

Tuning IOPS Performance: Mixed-Use Workload Micron 5100 SSD Example

When the SSD is fresh from the factory, its write performance is its highest. As applications write and rewrite to the SSD, write performance changes until it reaches a steady state, after which write performance remains relatively constant. For enterprise SSDs steady state is the primary performance state of interest.

This section describes an example of how we can use Micron Flex Capacity to tune the results of a Micron 5100 SATA SSD with a small-transfer (8K), random, mixed-use workload: 67% read and 33% write. Similar tuning applies to the Micron 9200 SSD with NVMe, but the values may differ to differences between the two drives.

We start with a 960GB (factory default capacity, designed for read-centric use) 5100 SSD as the example performance baseline. We then reset the 5100 drive to 800GB and 480GB capacity, and compare the results. After each capacity test, we restored the 5100 drive to a fresh-out-of-box (FOB) state, used Micron Flex Capacity features to reset its user capacity, and repeated the test with the new capacity.

Figure 3 shows relative IOPS performance compared to the 960GB baseline capacity. The horizontal axis shows time (from FOB at the start of testing on the left) and the vertical axis is IOPS performance relative to the baseline 960GB (shown as a percentage improvement relative to baseline 960GB IOPS). In Figure 3, a taller line means improved IOPS performance compared to the baseline 960GB drive.

As expected, write performance for the 960GB, 800GB, and 480GB all start at the same reference IOPS level (far left). As the drives are filled and refilled, the write IOPS decreases with time until reaching steady state.

Capacity at 800GB: Setting the capacity at 800GB shows an 8K IOPS improvement as expected. Starting at approximately a third of the way through the test, the 800GB capacity write IOPS show the beginnings of improvement with the observed maximum at steady state (end of test) reflecting a 26% 8K mixed-use IOPS improvement.

- **Capacity at 480GB:** The 480GB capacity created with the Micron Flex Capacity feature shows a more significant improvement over the 960GB default. Again, starting at the same IOPS value at FOB (far left), the 480GB 8K IOPS performance also begins to diverge a third of the way through the test, reaching its maximum improvement of about 40% at steady state (far right).




Figure 3: 8K Mixed-Use IOPS Performance vs Flex Capacity Feature

Notes:

- These are example capacities on a Micron 5100 SATA SSD only. As noted earlier, the Micron 9200 SSD with NVMe also supports the Micron Flex Capacity feature, enabling you to set the exact write IOPS performance and user capacity you need. (All possible capacities have not been tested, so user validation during qualification is required.)
- Combined average performance may not result in a linear relationship when reads and writes are mixed. We cannot always take 100% write workload IOPS performance and average it with a

33% write workload IOPS performance to determine a 67% write IOPS workload performance because servicing a WRITE command in steady state requires more internal bandwidth than READ commands. Although workload averaging may work in select cases, this is not a safe (universal) assumption.

- Similarly, capacity averaging may not result in a linear relationship. The performance of a 960GB Micron 5100 SATA SSD averaged with the performance of a 480GB 5100 setting may not accurately model the results of a 720GB 5100 setting. Using Micron Flex Capacity features enables you to verify specific performance by resetting the drive to 720GB.
- It is expected that the performance change as the SSD is written to (the shape of the WRITE IOPS versus time curve) may differ depending on the Micron Flex Capacity feature setting.
- The Micron 9200 SSD may behave differently than the example drive. Both SSDs support the Micron Flex Capacity feature, but they are very different products with different interfaces and are targeted for different uses. They will show different write IOPS performance as their capacity is adjusted.

Flex Capacity Feature and Drive Writes per Day

All SSDs wear as they are written; they have a finite write lifespan. An SSD's lifespan is rated in total bytes written (TBW). To make planning, design, deployment and support simpler, an SSD's lifespan is often expressed in how many times you can write the SSD's capacity (completely fill the drive) every day over a five-year warranty period, which is also referred to as the drive writes per day (DWPD). You may also see the term drive fills per day (DFPD), which means the same thing. The following section uses a Micron 5100 SATA SSD as an example to show how TBW and DWPD are related and how using the Micron Flex Capacity feature affects the DWPD value. As with write IOPS, the same principles apply to the Micron 9200 SSD with NVMe, but the calculated values will differ.

Micron Flex Capacity Feature and Changed DWPD

In this example, we use a Micron 9200 ECO SSD with NVMe, with an 11TB factory capacity setting and rated the endurance of 15.7PB over 5 years. Using these factory default values, we can calculate DWPD for five years:

$$\begin{aligned} \text{TBW} &= 15.7\text{PB} \\ \text{Time Period} &= 5 \text{ years} \\ \text{Factory capacity} &= 11\text{TB} \end{aligned}$$

For each day of the 5-year period, we can write

$$\begin{aligned} \text{GB/day} &= \text{Rated TBW} / \text{Time Period} \\ &= 15.7\text{PB} / 5 \text{ Years} \\ &= 16,076.8\text{TB} / 1826 \text{ days} \\ &= 8.8\text{TB} / \text{day} \end{aligned}$$

TBW is a fixed value; using the Micron Flex Capacity feature does not change the TBW. However, depending on the Micron Flex Capacity setting, the DWPD can change. Using the example default capacity of 11TB and the flex capacities of 8.8TB and 5.5TB, it is straightforward to calculate DWPD.

For any capacity, DWPD equals the amount of data that can be written per day based on its TBW value, divided by the drive capacity of interest:

$$\text{DWPD} = (\text{TB} / \text{day}) / 9200 \text{ capacity}$$

For the 11TB default capacity, take the TBW rating (145.7PB) and divide by the 5-year period, which yields about 8.8 TB/day — a value that we can use to calculate DWPD for three example 9200 capacities.

Starting with the factory default capacity of 11TB, we have:

$$\begin{aligned} \mathbf{11\text{TB}} \text{ DWPD} &= 8.8\text{TB (per day)} / \\ &11\text{TB capacity} \\ &= \mathbf{0.8 \text{ DWPD}} \end{aligned}$$

When we use the Micron Flex Capacity to resize this same 9200 drive to 8.8TB capacity, the DWPD changes:

$$\begin{aligned} \mathbf{8.8\text{TB}} \text{ DWPD} &= 8.8\text{TB (per day)} / \\ &8.8\text{TB capacity} \\ &= \mathbf{1.0 \text{ DWPD}} \end{aligned}$$

And finally, if we resize to 5.5TB:

$$\frac{5.5\text{TB DWPD} = 8.8\text{TB (per day)} / 5.5\text{TB capacity}}{5.5\text{TB capacity}} = 1.6 \text{ DWPD}$$

Although the TBW is the same, the tuned capacity changes the DWPD. In general, we expect DWPD to increase when we decrease the available capacity. Since the drive is available in three base platforms, the relative DWPD will have greater impact as you move from the ECO to the PRO to the MAX drive.

Similar calculations apply to the Mmicron 5100 SATA SSD. As we use the Micron Flex Capacity feature to reduce the drive's capacity, we change its effective DWPD. TBW does not change with changes to usable capacity.

Using the Micron Flex Capacity Feature With Micron's Storage Executive

Micron's Storage Executive tool unifies the method of capacity and performance tuning across all of our SSDs and is the preferred tool to use with the Micron Flex Capacity feature.

Storage Executive is available for download from micron.com, along with complete documentation. Implementing the Micron Flex Capacity feature via Storage Executive is beyond the scope of this paper. Please refer to the Storage Executive User Guide for specific instructions.

Conclusion

Different applications and workloads demand different storage for best results. With data center mainstream storage moving rapidly toward SSDs, precisely optimized configurations for both IOPS performance and available capacity are in greater demand.

With the Micron Flex Capacity feature, planners, designers, implementers, and administrators no longer have to compromise because of the limited number of SSD configurations, performance options, and capacities. The Micron Flex Capacity feature lets you easily create application-tunable SSDs in a flash.

Whether your applications and workloads need greater capacity with a read-focused workload that is highly cost-focused, greater write IOPS performance for write-intensive workloads, or performance to work well for mixed uses (with a read/write balance), the Flex Capacity feature enables the precise performance and capacity tuning you need while helping minimize qualification cycles and inventory (spares) and cost.

micron.com

Products are warranted only to meet Micron's production data sheet specifications. Products and specifications are subject to change without notice.
©2016 Micron Technology, Inc. Micron, the Micron logo, FlexPro and Flex Capacity are trademarks of Micron Technology, Inc. NVM Express is a registered trademark and NVMe is a trademark of NVM Express, Inc. All other trademarks are the property of their respective owners. All rights reserved. Rev. B 08/17, CCMMD-676576390-10578