

CONTENTS

SATA.....	1
ENTRY-LEVEL SATA.....	1
List of Entry-Level SATA drives.....	1
Typical Controllers in Entry-Level SATA Drives	2
MID-RANGE SATA	2
List of Mid-Range SATA Drives.....	3
Typical Controllers in Mid-Range SATA Drives	3
HIGH-END SATA	3
List of High-End SATA Drives.....	3
Typical Controllers in High-End SATA Drives	4
NVMe.....	5
ENTRY-LEVEL NVMe.....	5
List of Entry-Level NVMe Drives	5
Typical Controllers in Entry-Level NVMe Drives	6
MID-RANGE NVMe	6
List of Mid-Range NVMe Drives.....	7
HIGH-END NVMe	8
List of High-End NVMe Drives.....	8
Typical Controllers in High-End NVMe Drives.....	9
Glossary	11
Donators	18

SATA

SATA solid state drives (SSD) are commonplace. They utilize the SATA interface, whether in the 2.5" (physical) or M.2 (logical) form factor. They use the older AHCI protocol, a protocol not designed for solid state drives. In practice, any SSD will be an upgrade over a mechanical HDD and even old computers will have one or more SATA connectors.

One specific characteristic divides the "good" drives from the "bad": the presence or lack of DRAM. DRAM on a SSD does not act as a write or data cache as you would find on a HDD. Instead, it helps with the flash translation layer (FTL) – the method by which the drive communicates between the OS and the flash. The FTL handles myriad functions and the presence of DRAM assists it in these tasks, improving both performance and endurance. Drives utilized with mixed workloads – reads and writes – and especially many small operations benefit most from DRAM. This includes OS and mixed-use.

SATA SSDs in enclosures can act as fast portable devices. These enclosures rely on a protocol known as the USB Attached SCSI Protocol, or UASP, which improves performance over old methods and can pass the SCSI UNMAP command, an analogue for the SATA TRIM command. The TRIM command helps maintain performance and endurance by marking unused blocks. Going over USB has innate performance penalties and overhead – which will reduce performance, especially 4K and writes more than reads – but is somewhat mitigated by DRAM. For this reason, DRAM-equipped SATAs are ideal for current consoles when used externally, particularly since they do not properly support TRIM.

ENTRY-LEVEL SATA

Entry-Level SATA drives are suitable for light OS use or general data storage or may act as a dedicated games drive. Generally, it is desirable for these to be smaller in capacity for the former and larger in capacity for the latter, and they not be ideal as a primary (system) drive. For example, most QLC drives fall into this category because they require more flash to hit peak performance due to denser flash and the parallel nature of SSD design. QLC drives are also slower and less consistent outside of the temporary SLC cache. Entry-Level SATA drives also usually lack DRAM which is more crucial for SATA.

LIST OF ENTRY-LEVEL SATA DRIVES

[ADATA SU630/SU635](#)

[ADATA SU650/SU655](#)

[ADATA SU750/SU760](#)

[Crucial BX500](#)

HP S700/M700
Hyundai Sapphire
Inland SATA
Kingston A400/Q500
Lexar NS100
Mushkin Raw
Mushkin Source
Patriot Burst
Patriot P200
PNY CS900
PNY CS2311
Samsung 860/870 QVO
SanDisk SSD Plus
SP A55/S55
Team AX1/AX2
Team CX1/CX2
Team EX2
Team GX1/GX2
Team Vulcan G

TYPICAL CONTROLLERS IN ENTRY-LEVEL SATA DRIVES

Marvell 88NV1120
Maxio MAS0902A
Phison S11/T
Phison S13/T
Realtek RTS5733
SMI SM2258XT
SMI SM2259XT

MID-RANGE SATA

These SATA SSDs are suitable for any use but are a cheaper alternative to the very best drives on the market. They usually have DRAM but may be deficient in some other area: older or lower-quality flash, an older controller, etc. These are an option if budget is important but you do not want to drop to a DRAM-less drive.

LIST OF MID-RANGE SATA DRIVES

ADATA SU800/SX850
ADATA SX950U
Addlink S20
Gigabyte UD Pro
HP M700 (MLC)
HP S700 Pro
Hynix Gold S31
Kingston UV500
Mushkin Reactor (MLC)
PNY CS1311
Seagate BarraCuda
Team Delta RGB
Team L5 Lite 3D
Team Vulcan

TYPICAL CONTROLLERS IN MID-RANGE SATA DRIVES

Marvell 88SS1074
Phison S10
SMI SM2258

HIGH-END SATA

These are the very best SATA drives on the market. They have newer, powerful controllers and the newest flash. They are suitable for any type of usage.

LIST OF HIGH-END SATA DRIVES

Crucial MX500
Hynix Gold S31
Intel 545s
Kingston KC600
Lexar NS200
Samsung 860/870 EVO
SanDisk Ultra 3D

Seagate BarraCuda 120
WD Blue 3D

TYPICAL CONTROLLERS IN HIGH-END SATA DRIVES

Marvell 88SS1074

Phison S12

SMI SM2258

SMI SM2259

NVME

The non-volatile memory express (NVMe) protocol exists for solid state devices. It has many advantages over AHCI including new commands, deeper queue depth, improved interrupts, better parallelism (threads), better efficiency, lower I/O and processing overhead, higher IOPS, superior latency, higher bandwidth, etc. General users may only see minor performance gains, for example with game load times, with an NVMe drive when compared to SATA. Support is also reduced as these drives require a specialized M.2 socket and consumer boards have limited lanes and therefore bandwidth available.

However, support is becoming more common with, for example, AMD's X570 platform that also supports PCIe 4.0-based devices. PCIe drives saw more OEM orders in 2019 than SATA for the first time ever. Price premiums for NVMe are at an all-time low. Therefore, it is wise to consider an NVMe drive if your system supports it and your budget allows. These drives also have enclosures using different bridge controllers/chips but are generally limited in sequential performance by the interface.

Not all NVMe drives are equal. Mobile devices and small form factor (SFF) machines may put a premium on power efficiency, a smaller thermal footprint, or require a single-sided drive. Likewise, DRAM is not as critical on NVMe drives – the protocol has enough uplift over AHCI that it is optional for general use – and some support using system memory through host memory buffer (HMB) instead. We have QLC for capacity and TLC for consistency. Further, drives will have different controllers – four- or eight-channel, one or more cores, etc. – as well as different SLC caching designs. Therefore, differentiation is required.

ENTRY-LEVEL NVME

Entry-Level NVMe drives are entry-level or SATA replacement drives, or possibly ideal options for mobile and SFF machines. These drives tend to have weaker, four-channel controllers with less or no DRAM. Most of the DRAM-less options support HMB (note that HMB has no support with USB enclosure bridge chips). This includes both QLC- and TLC-based drives. These drives are also usually single-sided.

LIST OF ENTRY-LEVEL NVME DRIVES

[ADATA SX6000 Lite/Pro](#)
[ADATA Falcon](#)
[ADATA Swordfish](#)
[Crucial MP400](#)
[Crucial P1/P2](#)

Crucial P3/P3 Plus
Gigabyte NVMe
HP EX900
HP EX900 Plus
Inland Professional
Intel 660p/665p/670p
Kingston A1000/A2000
Kingston NV2
Lexar NM620
MDSSD SBX/SBXe
MSI Spatium M390
Mushkin Helix-L
Patriot Scorch
Plextor P300
Samsung 980 Non-PRO
Solidigm P41 Plus
SP P32A80
SP P34A60
Team MP32/MP33
WD SN350
WD SN500/SN550
WD SN570/SN580
WD SN750 SE

TYPICAL CONTROLLERS IN ENTRY-LEVEL NVME DRIVES

InnoGrit IG5208/IG5216
Realtek RTS5762DL/RTS5763DL
SMI SM2263/XT
Phison E8/E8T
Phison E13T
Phison E15T/E19T
Phison E21T (QLC)

MID-RANGE NVME

These are the best consumer drives with high performance on desktop. Typically found with eight-channel, powerful controllers, sufficient DRAM, solid flash, etc. Drives in this category will give you the best user experience while generally being affordable. Over time these may be phased out or replaced with newer hardware.

LIST OF MID-RANGE NVME DRIVES

ADATA ATOM 50/Legend 840

ADATA S40G

ADATA S50

ADATA SX8100/S7

ADATA SX8200/S11

ADATA SX8200/S11 Pro

ADATA S50 Lite (Gen4)

Addlink S70

Addlink S90 Lite

Corsair MP510

Corsair MP600

Crucial P5

Gigabyte Gen4 Aorus

HP EX920

HP EX950

HP FX900

Inland Performance

Inland Premium/Platinum

Intel 760p

Kingston KC2000/KC2500

Lexar NM790

MDSSD BPX Pro

Mushkin Delta

Mushkin Pilot

Mushkin Pilot-E

Patriot VP4100

Patriot VPN/VPR100

Plextor M9P+

PNY CS3030

PNY CS4040

Sabrent Rocket

Sabrent Rocket 4.0
Sabrent Rocket Q (2TB+)
Samsung 970 EVO Plus
Seagate Barra/Fire 510
Seagate FireCuda 520
Silicon Power US70
SK hynix P31
SP UD90
Team Cardea II/Liquid
Team Cardea Zero/Z440
Team MP34/Z340
Team MP44L
WD SN750
WD SN770

TYPICAL CONTROLLERS IN MID-RANGE NVME DRIVES

InnoGrit IG5220
Marvell 88SS1092
Phison E12/E12S
Phison E16
Phison E21T (TLC)
Realtek RTS5762
SMI SM2262/EN
SMI SM2267
SMI SM2269XT (TLC)

HIGH-END NVME

High-End NVMe SSDs are ones specialized for certain tasks, for example content creation or workstation workloads. These drives will have unique characteristics to aid in this type of usage, or may simply offer high-performance through a Gen4 interface and the newest flash. As such, these drives tend to be more expensive.

LIST OF HIGH-END NVME DRIVES

Acer GM7000

ADATA Legend 960/Max
ADATA S70/Blade
Aorus 10000 (Gen5)
Corsair MP600 Pro/Pro XT/LPX
Corsair MP700 (Gen5)
Crucial P5 Plus
Crucial T700 (Gen5)
Gigabyte Aorus 10/12K (Gen5)
HP FX900 Pro
Inland Performance Plus/GPP
Inland TD510 (Gen5)
Kingston KC3000/Fury
MSI 570/570 Pro (gen5)
Mushkin Gamma
Netac NV7000
Patriot VP4300
Plextor M10P(GN)
Sabrent Rocket 4.0 Plus/Plus-G
Sabrent Rocket X5 (Gen5)
Samsung 980/990 PRO
Seagate FireCuda 540 (Gen5)
SK hynix Plat. P41
Solidigm P44 Pro
Team Cardea A440/Pro/SE
Team Z540(Gen5)
Team Z54A (Gen5)
WD SN850/SN850X/SN850P

TYPICAL CONTROLLERS IN HIGH-END NVME DRIVES

InnoGrit IG5236
InnoGrit IG5666
Phison E18
Phison E26
SMI SM2264

GLOSSARY

2.5" – typical form factor for SATA, SATAe, and U.2 drives. Can be of varying height, usually 7mm.

2D/planar – two-dimensional NAND. Older type of NAND that is moving out of production except for budget-oriented drives. Also called planar because it sits on a plane.

3D – three-dimensional NAND. Newer type of NAND with a varying amount of layers. The third dimension offers more space between cells, reducing cell-to-cell interference. This allows greater headroom for a balance between higher density, higher performance, and endurance. This is generational, so a product listed as 96-layer maybe actually be 92-layer for example.

3D XPoint – type of memristor phase change memory/material (PCM).

32L – 32-layer 3D NAND.

64L – 64-layer 3D NAND.

96L – 96-layer 3D NAND.

128L – 128-layer 3D NAND.

176L – 176-layer 3D NAND.

Adapter – device to convert from one interface to another.

AHCI – advanced host controller interface. An old standard that did not anticipate the very low latencies and very high IOPS of solid state devices.

ARC – Argonaut RISC Core, a type of microcontroller used for SSD controllers along with ARM (see below).

ARM – Arm Holdings. Designer of microcontrollers including the ones found in SSDs.

BCH – Bose-Chaudhuri-Hocquenghem codes, an ECC scheme as used on some SSD controllers like the Phison S10.

BiCS – bit cost scaling; in general BiCS = 32-layer, BiCS2 = 48-layer, BiCS3 = 64-layer, BiCS4 = 96-layer, and BiCS5 = 128-layer. Based on the concept that 3D NAND is scalable and can drive down cost per capacity with more layers.

Bitline – string of cells.

Block – unit of storage in NAND containing many pages. Smallest unit of NAND available for erasure.

Block allocation table – table tracking the status of blocks.

Bridge chip/controller – usually a separate controller used to convert between interfaces, as found in enclosures.

Caching – collecting and combining writes and/or buffering accessed data in a faster storage device for a slower one.

CE – chip enable. Effectively, the maximum number of targets or dies per controller channel.

Compression – reducing the size of data.

Conversion – conversion of SLC blocks to native NAND (e.g. TLC) used with dynamic SLC cache, for example Micron's DWA.

Copyback – copying data within a plane.

DDR3 – double data rate 3rd generation SDRAM. DRAM used in older computers and as DRAM cache on many SSDs. Also found in a low-power configuration (LPDDR3).

DDR4 – double data rate 4th generation SDRAM. Type of DRAM now common in computers and for DRAM cache on some SSDs. Also found in a low-power configuration (LPDDR4).

Die – unit of NAND made up of blocks arranged into chips/packages of flash.

Direct-to-TLC/QLC – method of writing directly to native NAND when the SLC cache is exhausted.

DRAM – dynamic random-access memory. Volatile memory that is significantly faster than the flash used in a SSD, used as cache for mapping/translation and wear-leveling data.

DWA – dynamic write acceleration. Micron's method of converting blocks between SLC and the native format on-the-fly as deemed necessary. Also known as dynamic SLC.

DWPD – drive writes per day. This is a measure of possible writes to the NAND per day within the warrantied TBW.

ECC – error-correcting code. Used in SSDs to correct data errors, with limits. ECC is a critical aspect of SSD endurance. Some controllers, like the Phison S10, suffer for using BCH-based ECC in comparison to the now more-common LDPC-based ECC.

Encoding – overhead involved with the payload of data.

Endurance – amount of wear a drive might sustain while remaining operable. TBW and DWPD are for warranted writes, actual endurance is often far higher.

Folding – transfer of data from SLC blocks to a MLC/TLC/QLC block, for example three SLC blocks to one TLC block. Also a method of garbage collection whereby multiple partial blocks combine into one replacement block. Reduces NAND wear as it folds sequentially but is less performant. Proceeds on-die without intervention of the controller or DRAM in modern drives (copyback).

FTL – flash translation layer. Enables the OS to treat the SSD like a HDD, providing translation of addresses and garbage collection among other tasks. Processes often in DRAM and SLC cache for faster accesses.

GC – garbage collection. The consolidation of written data to allow for the erasure of blocks, usually in tandem with wear-leveling and TRIM.

HDD – hard disk drive. A mechanical/magnetic hard drive utilizing heads and platters for non-volatile storage.

HMB – host memory buffer. An NVMe technology that allows a SSD to utilize system memory (RAM) for DRAM cache instead of having dedicated memory on the drive's PCB.

IMFT – IM Flash Technologies. Joint-founded by Intel and Micron, recently fully acquired by Micron. Manufactures NAND.

Interleaving – performing I/O operations on multiple LUNs simultaneously.

IOPS – input/output operations per second. This is a measure of how many input and output operations completed per second.

JEDEC – JEDEC solid state technology association, a consortium of companies working on open standards for NAND.

JTAG – Joint Test Action Group, an industry standard for verifying and testing PCBs.

LPDC – low density parity codes. A type of ECC increasingly used by modern SSDs.

LUN – logical unit number, referring to the smallest quantity/grouping of something, e.g. a plane if discussing parallelization of flash.

M.2 – previously NGFF or next generation form factor. A socket that accepts a varying length of device, not limited to SSDs, with “keying” of the connector/fingers to determine the interface. Can act as direct PCIe and/or expose as USB and SATA.

Metadata – data about data, such as addressing/mapping information.

MLC – multi-level cell, often taken to indicate double-level cell. Twice the capacity of SLC but with only one-tenth the endurance, and less performance as well. Two-bit with four voltage states.

MT/s – megatransfers per second. Number of operations per second that transfer data, related to ONFI/Toggle with NAND.

NAND – not-and logic gate. Equivalent to an AND gate followed by a NOT gate.

NVMe – NVM Express, or non-volatile memory express. A protocol designed for the future of storage with superior capabilities to AHCI.

ONFi – Open NAND Interface Working Group. A group focusing on open standards for NAND.

Overprovisioning – reserving NAND for the controller's usage to assist in wear-leveling, garbage collection, performance (static SLC and FTL), and other uses. Increases endurance and often write performance. Modern controllers will use any free space, whether in a partition or not, for dynamic overprovisioning with the use of TRIM.

Package – one or more NAND dies together in a single structure on a SSD's PCB.

Page – row of bits in cells that act as a unit. Smallest unit of NAND for writing/programming. Often 16KB in size with modern TLC.

Page table – table tracking the status and addresses of pages.

PCB – printed circuit board. The portion of the drive that supports and connects other components.

PCIe – PCI Express, or peripheral component interconnect express. Serial connection acting as a switch to provide communication between devices and the CPU.

PCIe bifurcation – binary splitting of PCIe lanes.

P/E – program/erase cycles. NAND will face erasure before new writing begins. This value is therefore how many writes the cells can withstand before the data becomes corrupt and unreadable. This is a measurement of endurance. A 1TB drive of 1,000 P/E QLC NAND can withstand, for example, about 1PB of writes before accounting for the WAF.

Plane – part of a NAND die allowing for simultaneous operations. Typically two or four planes per die.

PLP – power loss protection.

PMIC – power management integrated circuit.

pMLC – pseudo-MLC. TLC NAND acting in MLC mode as a variation of SLC caching (see below), using two times the capacity.

pSLC – pseudo-SLC. Usually TLC NAND acting in SLC mode. Referred to as “SLC cache,” this uses three times the capacity for the same amount of data with the advantage of faster speeds. pSLC can be static or dynamic and some drives have both. Blocks in SLC mode can move to blocks of native NAND in a process known as folding. pSLC will fold to TLC in the background and, when full, newer controllers can write directly to the TLC portion of the drive. Static and dynamic modes.

pTLC – pseudo-TLC. Usually QLC NAND acting in TLC mode, usually in conjunction with a separate and primary pSLC mode.

QLC – quad-level cell. Four times the capacity of SLC but with only one-third the endurance of TLC, and less performant. Four-bit with sixteen voltage states.

Reference voltage – reference voltage utilized to determine the value of a cell’s bit. Also used as a verify voltage when programming a cell.

Sanitize – method of formatting a drive, deleting data and the mapping table.

SATA – serial advanced technology attachment. Old bus interface/standard for storage devices.

SATA Express – also known as SATAe, interface supporting both SATA and PCIe devices.

SCSI – small computer system interface. Set of standards for computers to communicate with peripheral devices via commands.

SDRAM – synchronous dynamic random-access memory.

Secure erase – method of formatting a drive, deleting the mapping table.

SED – self-encrypting drive.

SKU – stock keeping unit. Distinct type of item, used to connote a specific model and capacity.

SLC – single-layer cell. The least capacious but most performant type of flash with the highest endurance. Single-bit with two voltage states.

SMART – Self-Monitoring, Analysis and Reporting Technology. Method of tracking device health.

SMI – Silicon Motion. Designer of many SSD controllers.

SNIA – Storage Networking Industry Association, develop of storage standards.

Split-level cell – special type of semicircular memory cell split in half.

SRAM – static random-access memory. All SSD controllers have this type of local memory in small amounts. It can be used like DRAM cache, but is faster; basically CPU cache.

SSD – solid state drive. Device using circuits as memory to provide non-volatile storage.

String-stacking – stacking two or more decks of flash vertically to increase the wordline/layer count.

Sub-block – partial block.

Sub-page – partial or logical page.

Sub-plane – binary division of a plane.

TBW – total bytes written. Amount of actual data written to the NAND.

Tiering – separation of storage units based on performance.

Tiles – type of paired structure in IMFT flash with independent circuitry.

TLC – triple-level cell. Three times the capacity of SLC but with only one-third the endurance of MLC, and less performant than both. Three-bit with eight voltage states.

TRIM/UNMAP – commands to mark data blocks as no longer needed. TRIM as a SATA command marks blocks as free for rewriting which improves performance and endurance. UNMAP is a SCSI command that works similarly for block reclamation.

TSV – through silicon via, vertical connection through a wafer/die.

U.2 – also known as SFF-8639, interface similar to SATA Express with PCIe and SATA lanes.

UASP – USB attached SCSI protocol. Uses the SCSI command set for USB devices, including external SSDs.

USB – universal serial bus. Standard for connections and communication between computers and peripheral devices. Several different iterations with varying bandwidth.

WAF – write amplification factor. The proportion of host and NAND writes compared to just host; if the drive ends up doing 15GB of writes for 10GB of final data, this is a WAF of 1.5, a common value with consumer workloads. Impacts wear.

WD – Western Digital. Manufacturer of storage devices including SSDs. Owner of SanDisk.

Wear-leveling – the act of wearing the flash blocks evenly.

Wordline – row of cells written/read at a time.

Joel Magaña

jessiejames

Shawn (Younseok) Lee

Anton Skor

John Iburg

Alexandre Mecteau

James Ostmeyer

Conrad

Jebediah

Kennedy St John

Hari