Bargain Linux Build and Fast Friendly File Systems with AMD Ryzen and F2FS

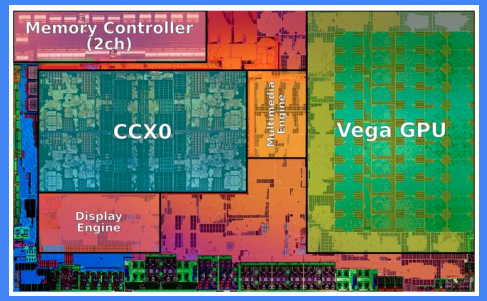

Wed, Aug 19, 2020

www.blu.org

#blu on freenode

https://meet.jit.si/blu.org

Brian DeLacey

https://en.wikichip.org/wiki/Category:microprocessor_models_by_amd_based_on_zen%2B

BLU Bargain Build

August 2020
Brian DeLacey

See Shankar’s Meeting from May

http://blu.org/cgi-bin/calendar/2020-may

An Overview of the Architecture of the Latest AMD Ryzen Chips
Date and Time

Wednesday, May 20, 2020 from 6:30 pm to 9:00 pm

Chat (Bill Ricker will be monitoring)

#blu on freenode

Applications - hexchat (join #blu),

Youtube Chat / Jitsi Chat / freenode

Top Chat vs. Livechat on Youtube

August 19 Meeting Agenda

In this three part meeting, we'll review the build list of components and overall budget for a new Linux-powered system with the AMD Ryzen 3. In the second part, we'll push this system to its byte-busting limits for powering a permanent archival system. In the third part, we'll talk about different on disk file formats and recording media types - and ask what matters more?

We'll begin with an overview of the build - covering the purchase of parts, continuing with assembly, and concluding with the installation of Ubuntu. Then we'll move to various considerations and options for the Terabyte Transportation Generation. We'll discuss managing backups, and showcase the capabilities of rclone and what you can do with affordable terabytes of storage.

We'll then look at different file systems and compare considerations for use in backups. We'll wrap it all up with a forensic review of the Flash-Friendly File System (F2FS) - discussing details of its design, data layout and aspects of encryption. We'll looking into the pros and cons of using F2FS as the underlying file system for a permanent archive of all your most valuable data. Finally, we'll discuss news reports of an SMR “scandal” and compare that technology to CMR.

Agenda

1. Build Overview / Parts and Labor
    a. Cost, supply chain
2. Application Usage / Consolidating Backups
    a. Rclone
3. F2FS / Flash-Friendly File System
    a. Design for Flash, useful for SMR?

Key Parts, Prices

Your store » Micro Center Cambridge
730 Memorial Drive
Cambridge , MA 02139
General Manager Arsen Askaryan
Text (617) 234-6400
www.microcenter.com

Your Sale Information

SKU	Description	Quantity	Price Per	Total Price
673996	COOLMAST N200 MATX MINI-TOWER CASE	1	49.99	49.99
951897	AMD AMD RYZEN 3 3200G WRAITHS S/N: 9HH2988N00002	1	79.99	79.99
402404	IPSG 650W 80+ SEMI ATX PSU S/N: 1939080003591650BR21F09016337	1	69.99	69.99
410936	IPSG 120GB I PRO SSD DRIVE	1	19.99	19.99
348201	G.SKILL 16GB 2X8GB DDR4 3200 KIT	1	69.99	69.99
862318	SEAGATE 2TB BARRACUDA 3.5" HD S/N: 763649123065	1	51.99	51.99
802538	ASUS PRIME B450M-A/CSM MATX S/N: K6M0KS033121LUK	1	74.99	74.99
	Subtotal »			$416.93
	Tax »			$26.06
	Sale TOTAL »			**$442.99**

Ryzen Facts

Facts about "Ryzen 3 3200G - AMD" RDF feed

Has subobject	*Ryzen 3 3200G - AMD#io* +
base frequency	3,600 MHz (3.6 GHz, 3,600,000 kHz) +
bus type	PCIe 3.0 +
clock multiplier	36 +
core count	4 +
core name	Picasso +
designer	AMD +
die area	209.78 mm² (0.325 in², 2.098 cm², 209,780,000 µm²) +
family	Ryzen 3 +
first announced	June 12, 2019 +
first launched	July 7, 2019 +
full page name	amd/ryzen 3/3200g +
has advanced vector extensions	true +
has advanced vector extensions 2	true +
has amd amd-v technology	true +
has amd amd-vi technology	true +
has amd sensemi technology	true +
has ecc memory support	false +
has feature	Advanced Vector Extensions +, Advanced Vector Extensions 2 +, Advanced Encryption Standard Instruction Set Extension + and SenseMI Technology +
has simultaneous multithreading	true +
has x86 advanced encryption standard instruction set extension	true +
instance of	microprocessor +
integrated gpu	Radeon Vega 8 +
integrated gpu designer	AMD +
integrated gpu max frequency	1,250 MHz (1.25 GHz, 1,250,000 KHz) +
isa	x86-64 +
isa family	x86 +

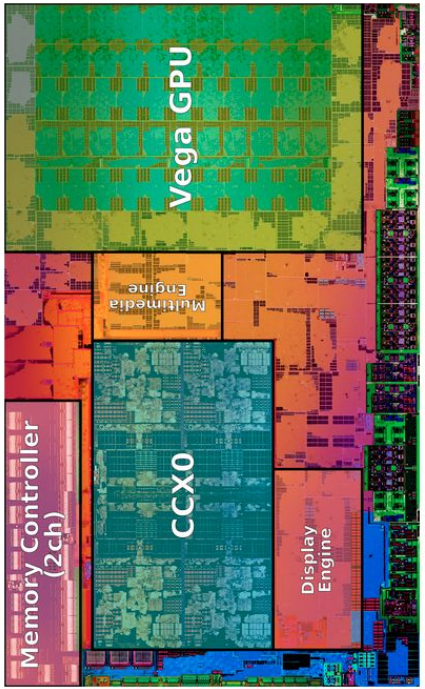

https://en.wikichip.org/wiki/amd/ryzen_3/3200g

l1$ size	0.375 MiB (384 KiB, 393,216 B, 3.662109e-4 GiB) +
l1d$ description	8-way set associative +
l1d$ size	0.125 MiB (128 KiB, 131,072 B, 1.220703e-4 GiB) +
l1i$ description	4-way set associative +
l1i$ size	0.25 MiB (256 KiB, 262,144 B, 2.441406e-4 GiB) +
l2$ description	8-way set associative +
l2$ size	2 MiB (2,048 KiB, 2,097,152 B, 0.00195 GiB) +
l3$ size	4 MiB (4,096 KiB, 4,194,304 B, 0.00391 GiB) +
ldate	July 7, 2019 +
manufacturer	GlobalFoundries +
market segment	Desktop +
max cpu count	1 +
max memory	65,536 MiB (67,108,864 KiB, 68,719,476,736 B, 64 GiB, 0.0625 TiB) +
max memory bandwidth	43.71 GiB/s (76.055 GB/s, 44,759.04 MiB/s, 0.0427 TiB/s, 0.0469 TB/s) +
max memory channels	2 +
max operating temperature	95 °C +
max pcie lanes	20 +
microarchitecture	Zen+ +
min operating temperature	0° C +
model number	3200G +
name	Ryzen 3 3200G +
package	µOPGA-1331 +
part number	YD3200C5M4MFH + and YD3200C5FHBOX +
process	12 nm (0.012 µm, 1.2e-5 mm) +
series	3000 +
smp max ways	1 +
socket	Socket AM4 +
supported memory type	DDR4-2933 +
tdp	65 W (65,000 mW, 0.0872 hp, 0.065 kW) +
tdp down	45 W (45,000 mW, 0.0603 hp, 0.045 kW) +
technology	CMOS +
thread count	4 +
transistor count	4,940,000,000 +
turbo frequency	4,000 MHz (4 GHz, 4,000,000 kHz) +
word size	64 bit (8 octets, 16 nibbles) +

Benchmarks

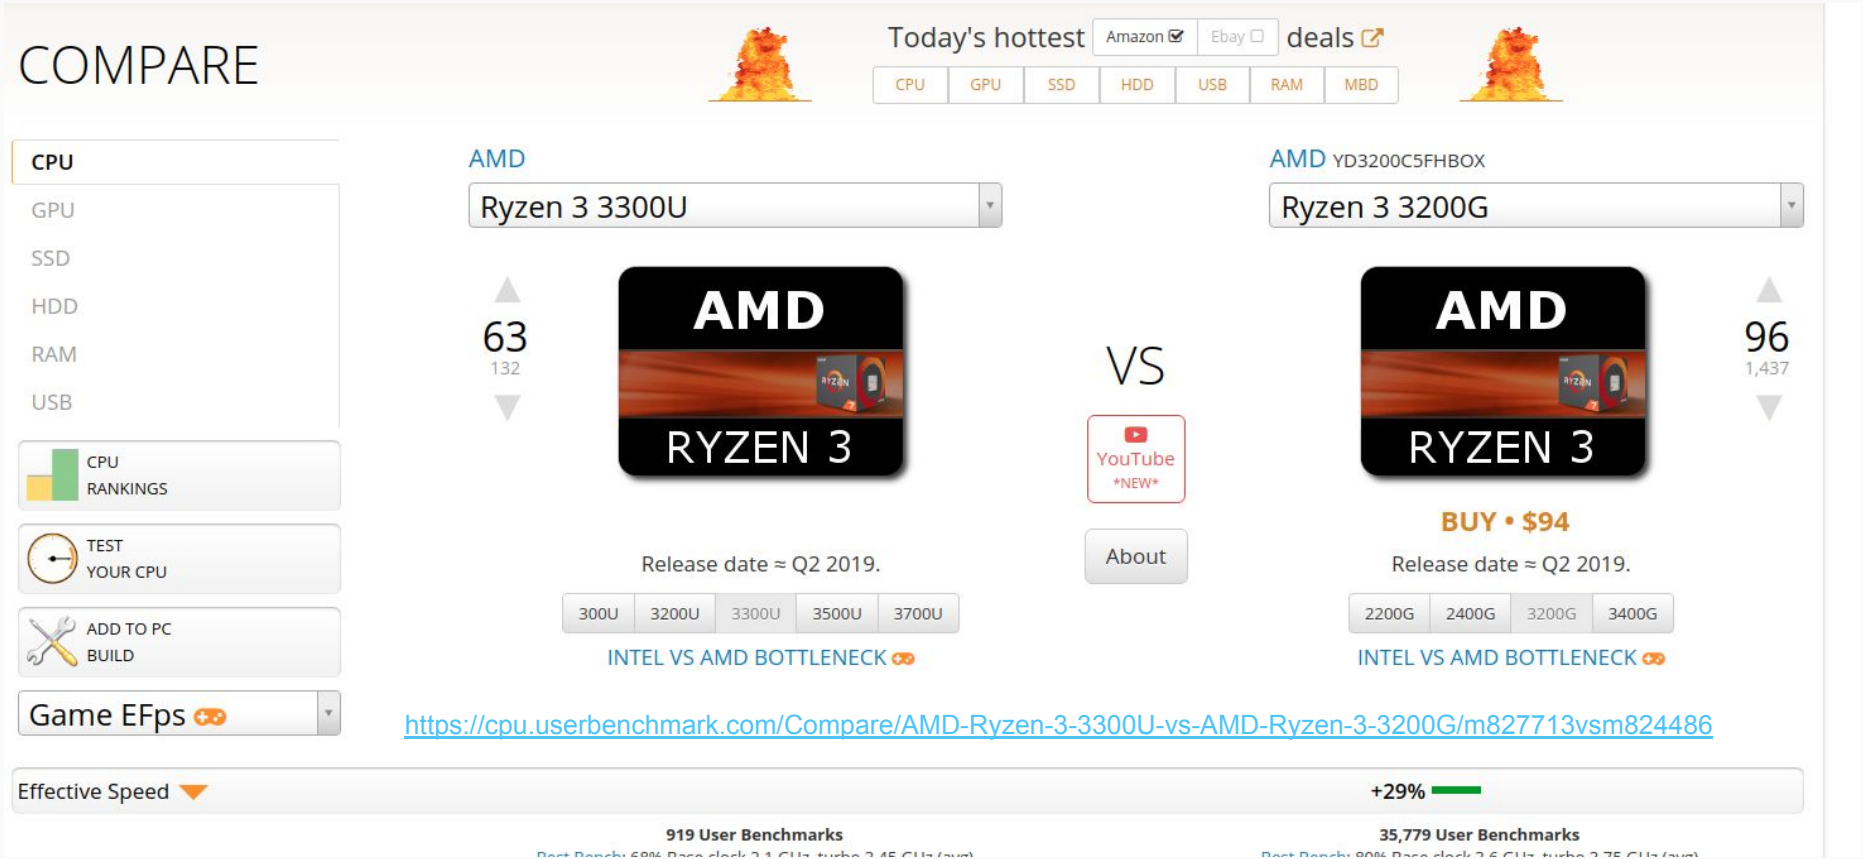
COMPARE
Today's hottest
Amazon
Ebay
deals
CPU
GPU
SSD
HDD
USB
RAM
MBD
CPU
GPU
SSD
HDD
RAM
USB
CPU RANKINGS
TEST YOUR CPU
ADD TO PC BUILD
Game EFps
AMD
Ryzen 3 3300U
AMD YD3200C5FHBOX
Ryzen 3 3200G
63
132
96
1,437
AMD
RYZEN 3
VS
YouTube
NEW
About
BUY • $94
Release date ≈ Q2 2019.
300U
3200U
3300U
3500U
3700U
2200G
2400G
3200G
3400G
INTEL VS AMD BOTTLENECK
INTEL VS AMD BOTTLENECK
https://cpu.userbenchmark.com/Compare/AMD-Ryzen-3-3300U-vs-AMD-Ryzen-3-3200G/m827713vsm824486
Effective Speed
+29%
919 User Benchmarks
35,779 User Benchmarks

Motile M141

The $199 Motile M141 With AMD Ryzen 3 3200U Offers Surprisingly Decent Performance: *Last week we published benchmarks of the Motile M141, Walmart's private-label tech branch, and the M141 being a Ryzen 3 3200U powered laptop that has been retailing for just $199 USD. In those initial benchmarks was an extensive look at the Windows vs. Linux performance while this article today is looking at the performance of this AMD Ryzen 3 laptop against a number of old and new Intel laptops, all tested using a daily snapshot of Ubuntu 20.04 LTS.*

Phronix.org, by Michael Larabel in Computers on 3 February 2020

"Eight laptops I had available were tested for putting the performance of this $199 USD laptop in perspective. Though as one unfortunate item: since running the original article and all the publicity on the Motile M141, Walmart has increased its price at least temporarily to $279 USD. We'll see if it falls back to $199 in the days ahead but even at $279 is still a decent deal."

M141-BK (as of 8/17/2020)

Motile

MOTILE 14" Performance Laptop, FHD, AMD Ryzen 3 with Radeon Vega 3 Graphics, THX Spatial Audio, Tuned by THX display, 4GB RAM, 128GB SSD, HDMI, Front 720P HD IR Camera - Black

Model: M141-BK Walmart # 577035574

★★★★☆ (4.0) 232 ratings Write a review

$449.00 ~~$599.00~~

$42/mo with affirm Learn how

Actual Color: Black

Ryzen Clocks

Model	Cores /Threads	Base Clock	Boost Clock	Cache	Graphics	iGPU Base Clock	TDP
Ryzen 5 3400G	**4 / 8**	**3.7 GHz**	**4.2 GHz**	**6MB**	**Radeon RX Vega 11**	**1400 MHz**	**65W**
Ryzen 5 2400G	4 / 8	3.6 GHz	3.9 GHz	4MB	Radeon RX Vega 11	1250 MHz	65W
Ryzen 3 3200G	**4 / 4**	**3.6 GHz**	**4.0 GHz**	**6MB**	**Radeon Vega 8**	**1250 MHz**	**65W**
Ryzen 3 2200G	4 / 4	3.5 GHz	3.7 GHz	4MB	Radeon Vega 8	1100 MHz	65W

https://www.tomshardware.com/news/amd-ryzen-3-3200g-ryzen-5-3400g-specs-pricing,39619.html

Boxes of Parts

Cool Complexity

Assembly

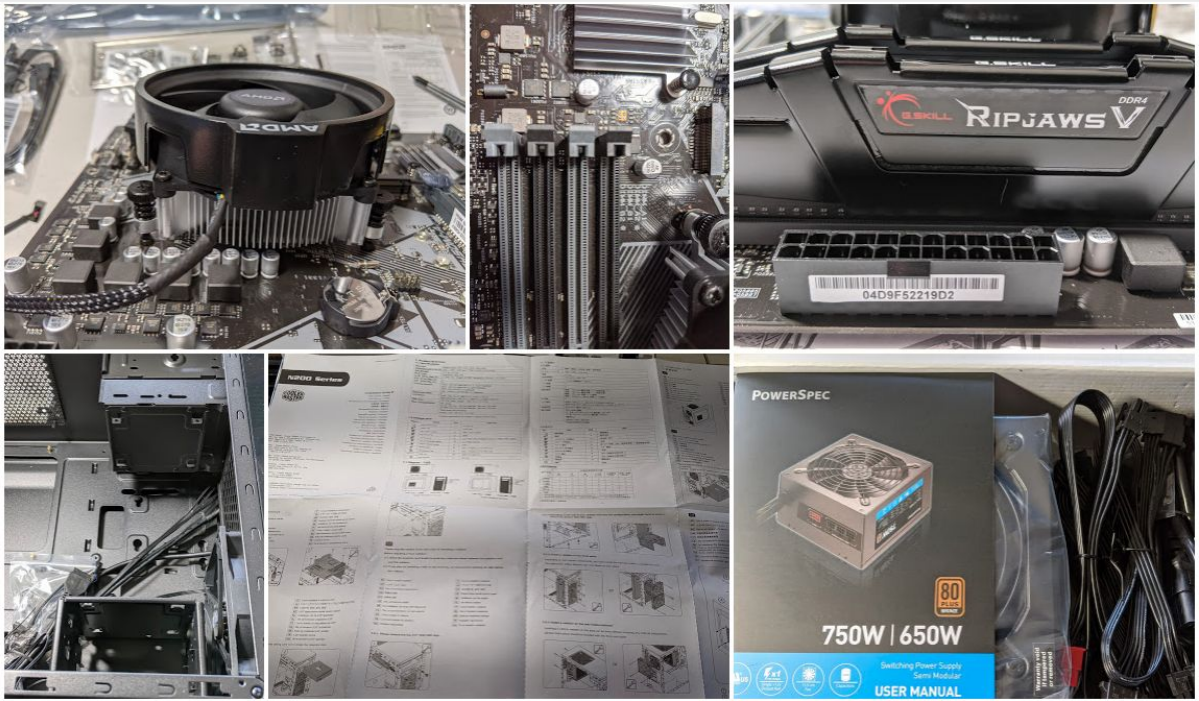
AMD
G.SKILL RIPJAWS V DDR4
04D9F52219D2
PowerSpec
80 PLUS
750W | 650W
Switching Power Supply
Semi Modular
USER MANUAL

In case you need to make connections ...

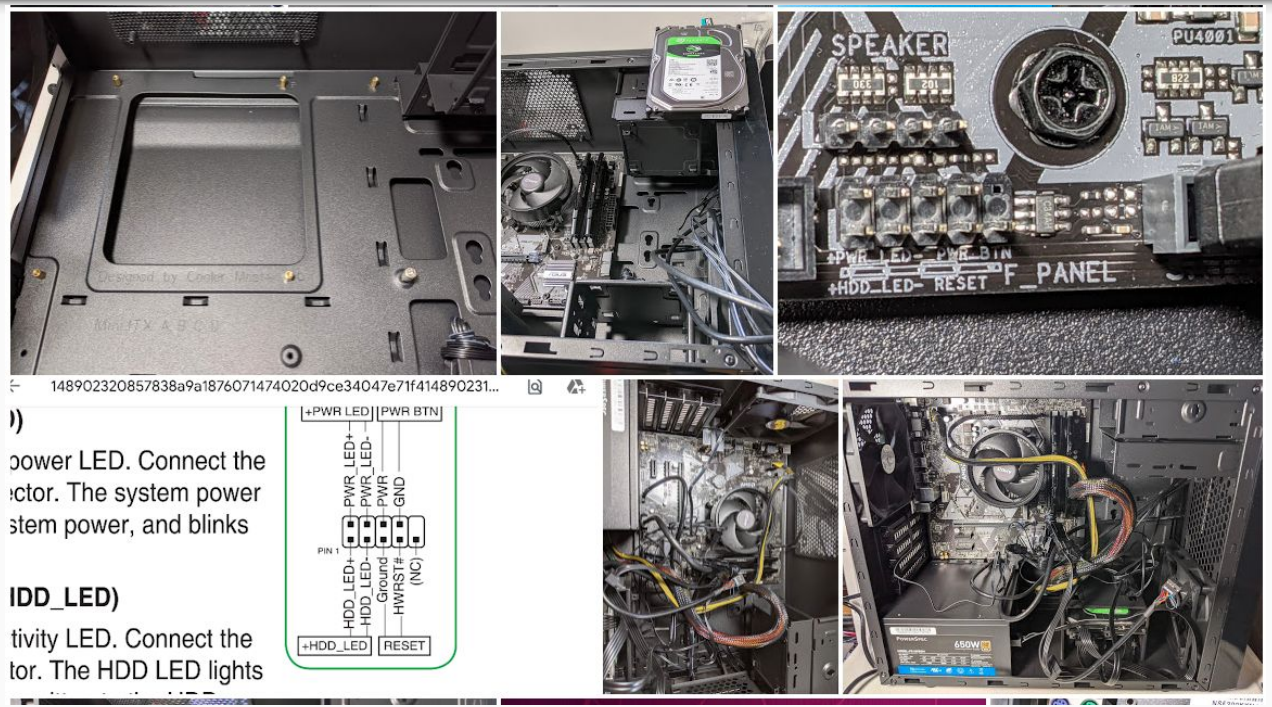

Performance Comp via Hardware Info

Intel Core i5, 8th Generation - i5-8250U @ 1.60 Ghz, 4 core with fpu

1. CPU Blowfish 1.27
2. CPU CryptoHash 813.99
3. CPU Fibonacci 0.55
4. CPU N-Queens 5.92
5. CPU Zlib 1.16
6. FPU FFT 0.89
7. FPU Raytracing 1.36
8. GPU Drawing 5283.12

AMD Ryzen 3 3200G with Radeon Vega Graphic, 4 core with fpu

1. CPU Blowfish 2.61 (Lower Better)
2. CPU CryptoHash 668.39 (Higher Better)
3. CPU Fibonacci 0.59 (Lower Better)
4. CPU N-Queens 3.90 (Lower Better)
5. CPU Zlib 1.17 (Higher Better)
6. FPU FFT 0.99 (Lower Better)
7. FPU Raytracing 1.75 (Lower Better)
8. GPU Drawing 6511.60 (Higher Better)

sudo apt-get install hardinfo

Consolidating Backups

Rclone copy

2.650 TB of data => 409,686 files

sudo mount /dev/sdd1 '/media/brian/Flex' -o noatime (this is USB2)

brian@B-AMD:~$ rclone copy '/media/brian/FreeAgent GoFlex Drive/bkup20160220' '/media/brian/Sea8T-X/bkup20160220-noatime' --stats-log-level NOTICE --skip-links

```
2020/07/20 05:05:46 NOTICE:
Transferred:        2.650T / 2.650 TBytes, 100%, 50.910 MBytes/s, ETA 0s
Errors:                 0
Checks:                 0 / 0, -
Transferred:       386066 / 386066, 100%
Elapsed time:  15h9m36.8s
```

Data Transportation

2.65 TB
409,686 files
15h9m36.8s
USB2 -> USB3
via rclone copy

In this 2.650 TB of data there are 409,686 files

==========

7/20/20 finished

sudo mount /dev/sdd1 '/media/brian/Flex' -o noatime (this is USB2)

brian@B-AMD:~$ rclone copy '/media/brian/FreeAgent GoFlex Drive/bkup20160220' '/media/brian/Sea8T-X/bkup20160220-noatime' --stats-log-level NOTICE --skip-links

```
2020/07/20 05:05:46 NOTICE:
Transferred:        2.650T / 2.650 TBytes, 100%, 50.910 MBytes/s, ETA 0s
Errors:                 0
Checks:                 0 / 0, -
Transferred:       386066 / 386066, 100%
Elapsed time:  15h9m36.8s
```

2020/07/20 05:05:46 NOTICE:
Transferred: 2.650T / 2.650 TBytes, 100%, 50.910 MBytes/s, ETA 0s
Errors: 0
Checks: 0 / 0, -
Transferred: 386066 / 386066, 100%
Elapsed time: 15h9m36.8s

How Fast are Drives?

External drive		Internal drive	
USB 1.1	Up to 12 Mbits/sec	UltraATA 100	Up to 100 Mbytes/sec
USB 2.0	Up to 480 Mbits/sec	Serial ATA 1.5	Up to 1.5 Gbits/sec
USB 3.0	Up to 4.8 Gbits/sec	SATA 3.0	Up to 3.0 Gbits/sec
		SATA 6.0	Up to 6.0 Gbits/sec
1394a (Firewire 400)	Up to 400 Mbits/sec		
1394b (Firewire 800)	Up to 800 Mbits/sec	Serial-Attached SCSI (SAS)	Up to 1.5, 3.0, 6.0, or 12.0 Gbits/sec
eSATA	Up to 1.5 or 3.0 Gbits/sec		
Thunderbolt	Up to 10 Gbits/sec		

Source: https://www.seagate.com/support/kb/how-fast-should-an-external-drive-be-172213en/

F2FS: A New File System for Flash Storage (2015)

Who uses F2FS?

“Motorola Mobility has used F2FS in their Moto G/E/X and Droid phones since 2012. Google first used F2FS in their Nexus 9 in 2014.[17] However Google's other products didn't adopt F2FS until the Pixel 3 when F2FS was updated with inline crypto hardware support.[18]”

https://en.wikipedia.org/wiki/F2FS

F2FS: The Flash Friendly File System

f2fs: introduce flash-friendly file system

https://lwn.net/Articles/518718/

If you'd like to experience f2fs, simply:

sudo apt install f2fs-tools

```
# mkfs.f2fs /dev/sdb1

# mount -t f2fs /dev/sdb1 /mnt/f2fs
```

Why F2FS?

“frequent random writes to an SSD would incur internal fragmentation of the underlying media and degrade the sustained SSD performance ...Unless handled carefully, frequent random writes and flush operations in modern workloads can seriously increase a flash device’s I/O latency and reduce the device lifetime.”

Adaptive logging
“F2FS builds basically on append-only logging to turn random writes into sequential ones. At high storage utilization, however, it changes the logging strategy to threaded logging [23] to avoid long write latency. In essence, threaded logging writes new data to free space in a dirty segment without cleaning it in the foreground. This strategy works well on modern flash devices but may not do so on HDDs.”

https://www.usenix.org/system/files/conference/fast15/fast15-paper-lee.pdf

F2FS: A New File System for Flash Storage Changman Lee, Dongho Sim, Joo-Young Hwang, and Sangyeun Cho, Samsung Electronics Co., Ltd.

https://www.usenix.org/conference/fast15/technical-sessions/presentation/lee

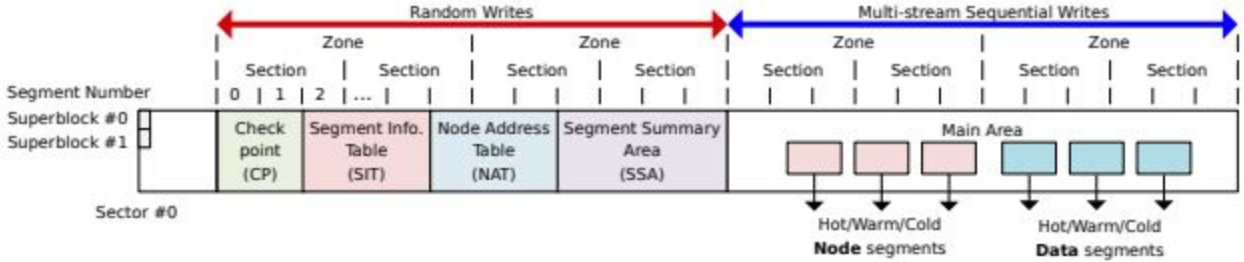

Figure 1: On-disk layout of F2FS.

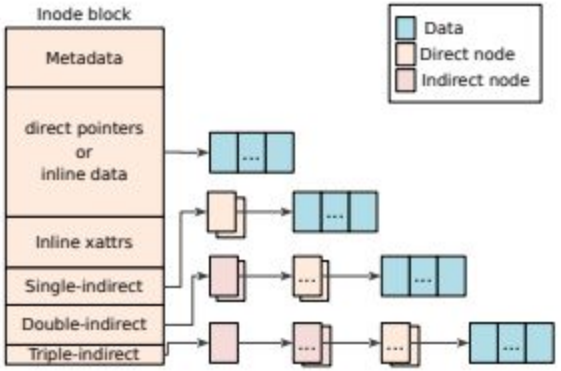

Figure 2: File structure of F2FS.

https://www.usenix.org/system/files/conference/fast15/fast15-paper-lee.pdf

“F2FS is a Linux file system designed to perform well on modern flash storage devices. The file system builds on append-only logging and its key design decisions were made with the characteristics of flash storage in mind. This paper describes the main design ideas, data structures, algorithms and the resulting performance of F2FS.

Experimental results highlight the desirable performance of F2FS; on a state-of-the-art mobile system, it outperforms EXT4 under synthetic workloads …”

https://www.usenix.org/conference/fast15/technical-sessions/presentation/lee

Creating the first F2FS

```
brian@hplaptop:~$ sudo mkfs.f2fs /dev/sdb

        F2FS-tools: mkfs.f2fs Ver: 1.11.0 (2018-07-10)

Info: Disable heap-based policy
Info: Debug level = 0
Info: Label =
Info: Trim is enabled
        /dev/sdb appears to contain a partition table (dos).
        Use the -f option to force overwrite.
brian@hplaptop:~$ sudo mkfs.f2fs /dev/sdb -f

        F2FS-tools: mkfs.f2fs Ver: 1.11.0 (2018-07-10)

Info: Disable heap-based policy
Info: Debug level = 0
Info: Label =
Info: Trim is enabled
Info: [/dev/sdb] Disk Model: FreeAgent GoFlex0148P!
                                                  ◆FreeAgent GoFlex
Info: Segments per section = 1
Info: Sections per zone = 1
Info: sector size = 512
Info: total sectors = 976773167 (476940 MB)
Info: zone aligned segment0 blkaddr: 512
Info: format version with
  "Linux version 5.4.0-37-generic (buildd@lcy01-amd64-001) (gcc version 9.3.0 (Ubuntu 9.3.0-10ubuntu
2)) #41-Ubuntu SMP Wed Jun 3 18:57:02 UTC 2020"
Info: [/dev/sdb] Discarding device
Info: This device doesn't support BLKSECDISCARD
Info: This device doesn't support BLKDISCARD
Info: Overprovision ratio = 0.290%
Info: Overprovision segments = 1384 (GC reserved = 697)
Info: format successful
```

F2FS and Enryption

Since Linux 4.2, **F2FS** natively supports file **encryption**. **Encryption** is applied at the directory level, and different directories can use different **encryption** keys. This is different from both dm-crypt, which is block-device level **encryption**, and from eCryptfs, which is a stacked cryptographic filesystem.Aug 11, 2020

https://wiki.archlinux.org/

Filesystem-level encryption (fscrypt)

Introduction

fscrypt is a library which filesystems can hook into to support transparent encryption of files and directories.

Note: “fscrypt” in this document refers to the kernel-level portion, implemented in `fs/crypto/`, as opposed to the userspace tool fscrypt. This document only covers the kernel-level portion. For command-line examples of how to use encryption, see the documentation for the userspace tool fscrypt. Also, it is recommended to use the fscrypt userspace tool, or other existing userspace tools such as fscryptctl or Android’s key management system, over using the kernel’s API directly. Using existing tools reduces the chance of introducing your own security bugs. (Nevertheless, for completeness this documentation covers the kernel’s API anyway.)

GParted Partition Look

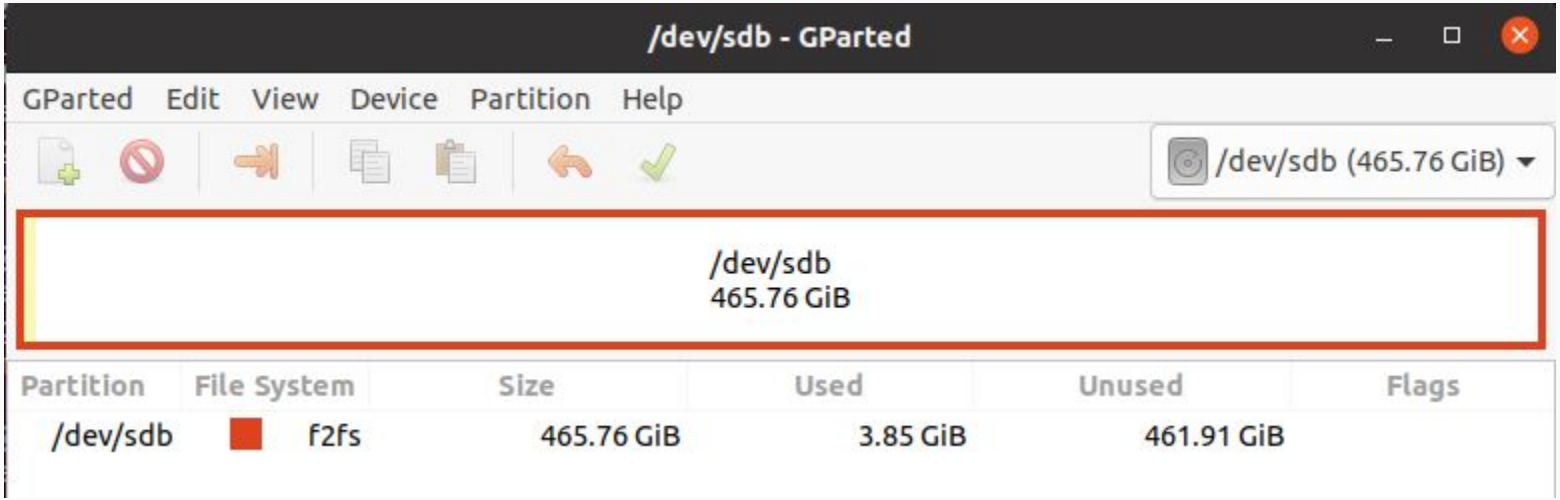

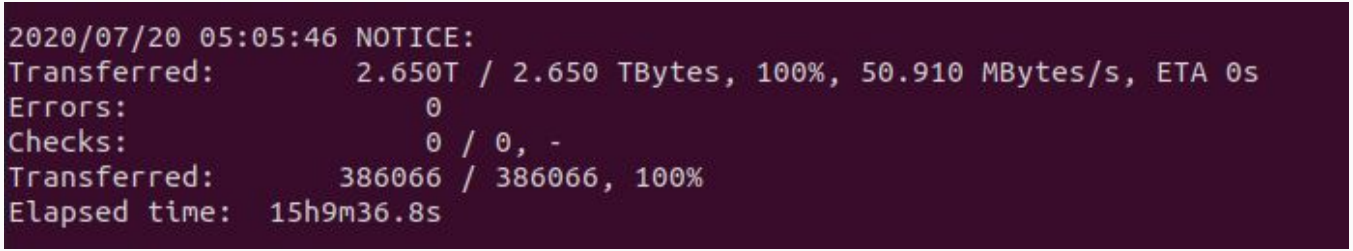

2.650 TB of data there are 409,686 files

sudo mount /dev/sdd1 '/media/brian/Flex' -o noatime (this is USB2)

brian@B-AMD:~$ rclone copy '/media/brian/FreeAgent GoFlex Drive/bkup20160220' '/media/brian/Sea8T-X/bkup20160220-noatime' --stats-log-level NOTICE --skip-links

Phronix Tests with F2FS (28 December 2018)

Written by Michael Larabel in Computers on 28 December 2018. **Page 2 of 3**. 17 Comments

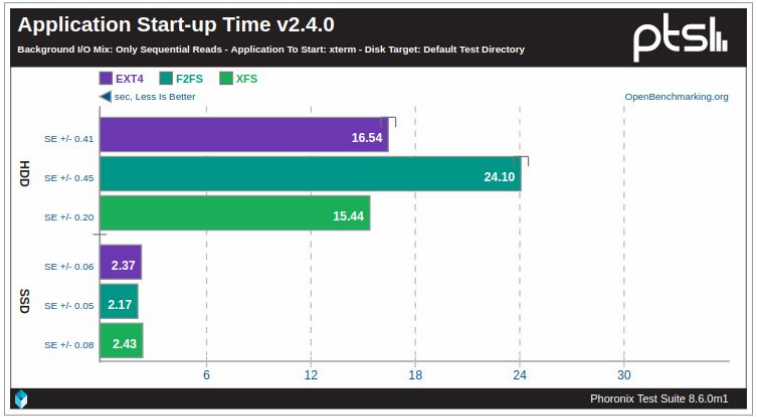

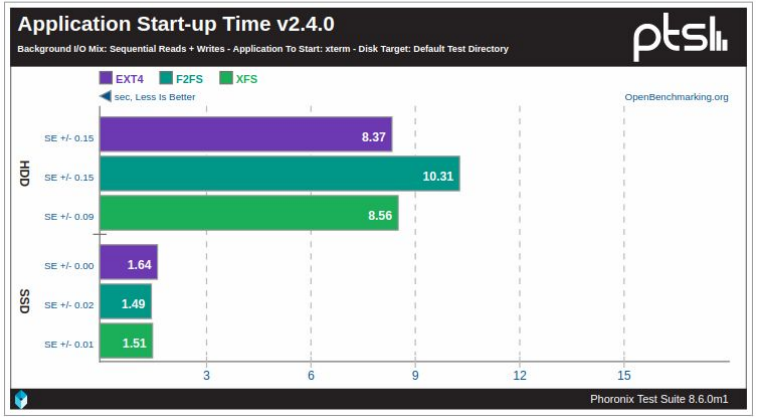

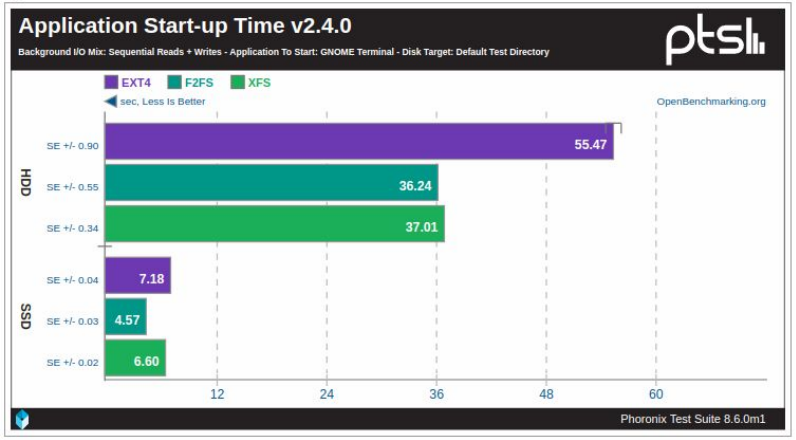

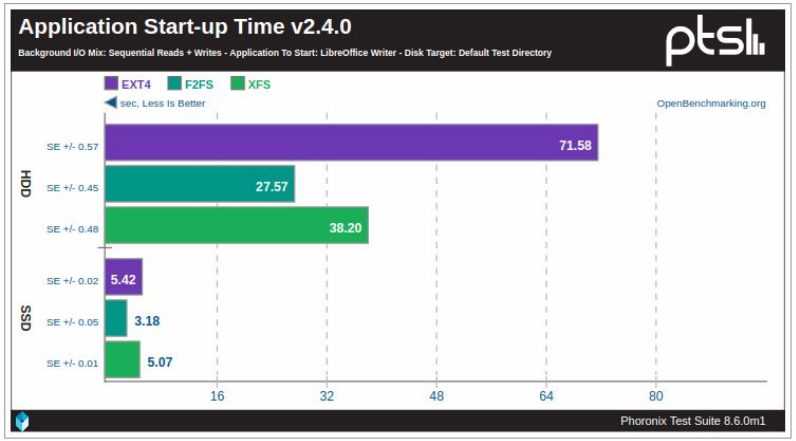

https://www.phoronix.com/scan.php?page=article&item=f2fs-hdd-test&num=2

Files (before mounting)

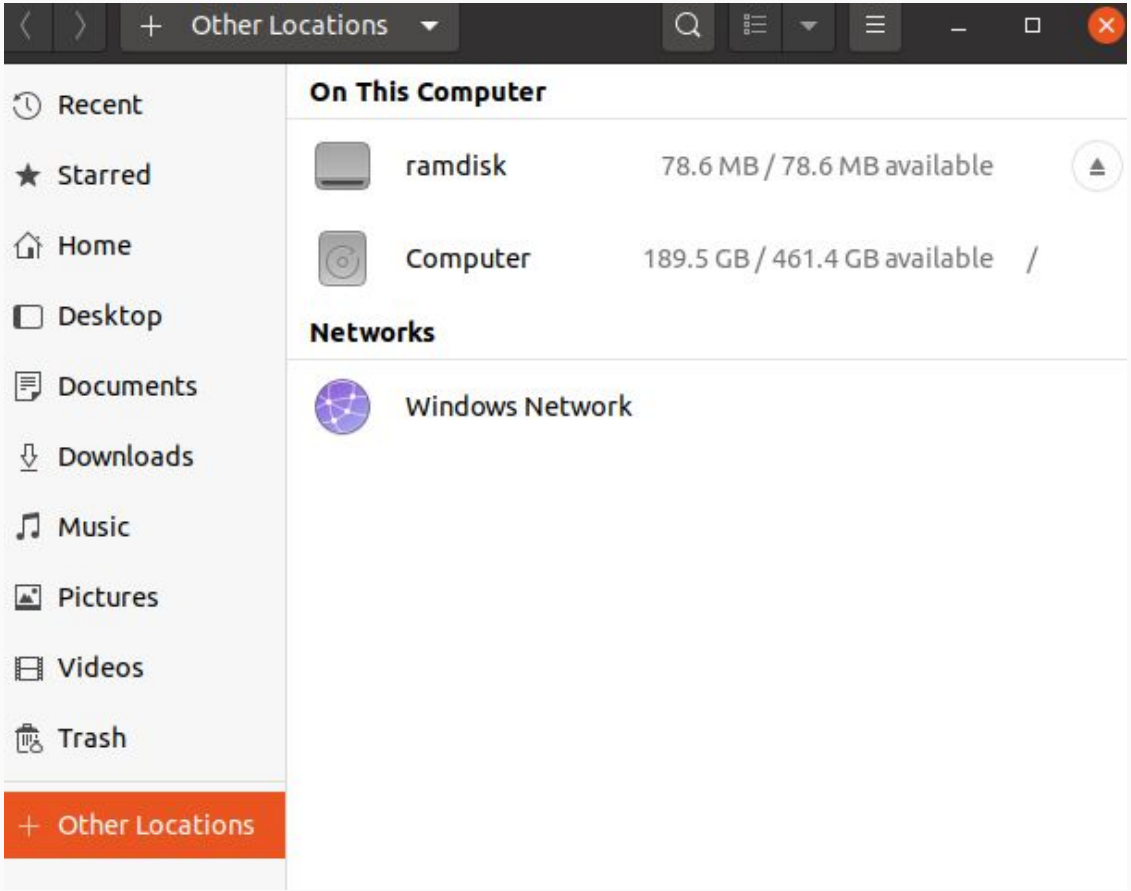

Phronix File System Testing

With the initial create process of Compile Bench, F2FS returned to being the fastest followed by EXT4 and then XFS.

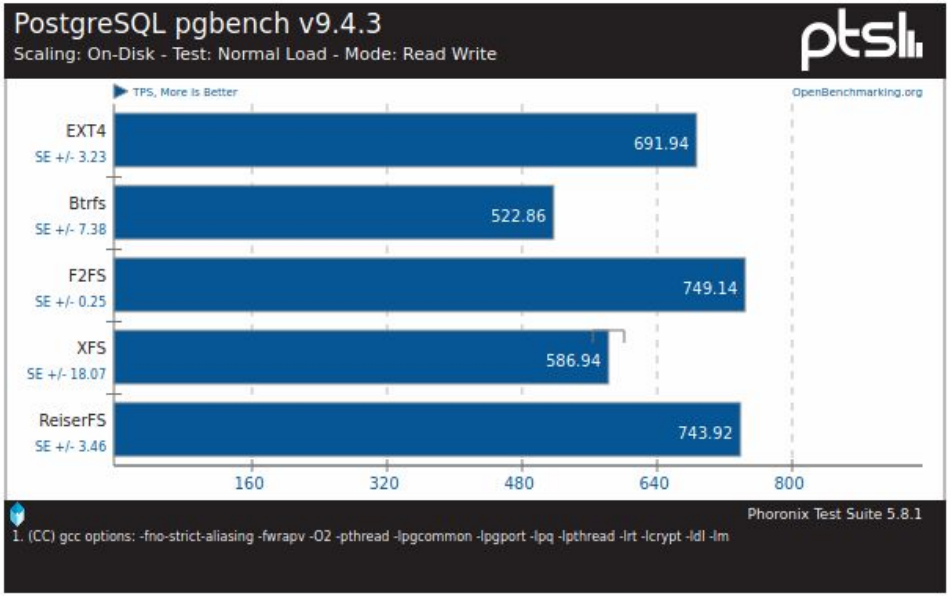

“Compilebench tries to age a filesystem by simulating some of the disk IO common in creating, compiling, patching, stating and reading kernel trees. It indirectly measures how well filesystems can maintain directory locality as the disk fills up and directories age. This current test is setup to use the makej mode with 10 initial directories”

Source:
https://openbenchmarking.org/test/pts/compilebench

https://www.phoronix.com/scan.php?page=article&item=linux-41-filesystem&num=3

Phronix pgbench

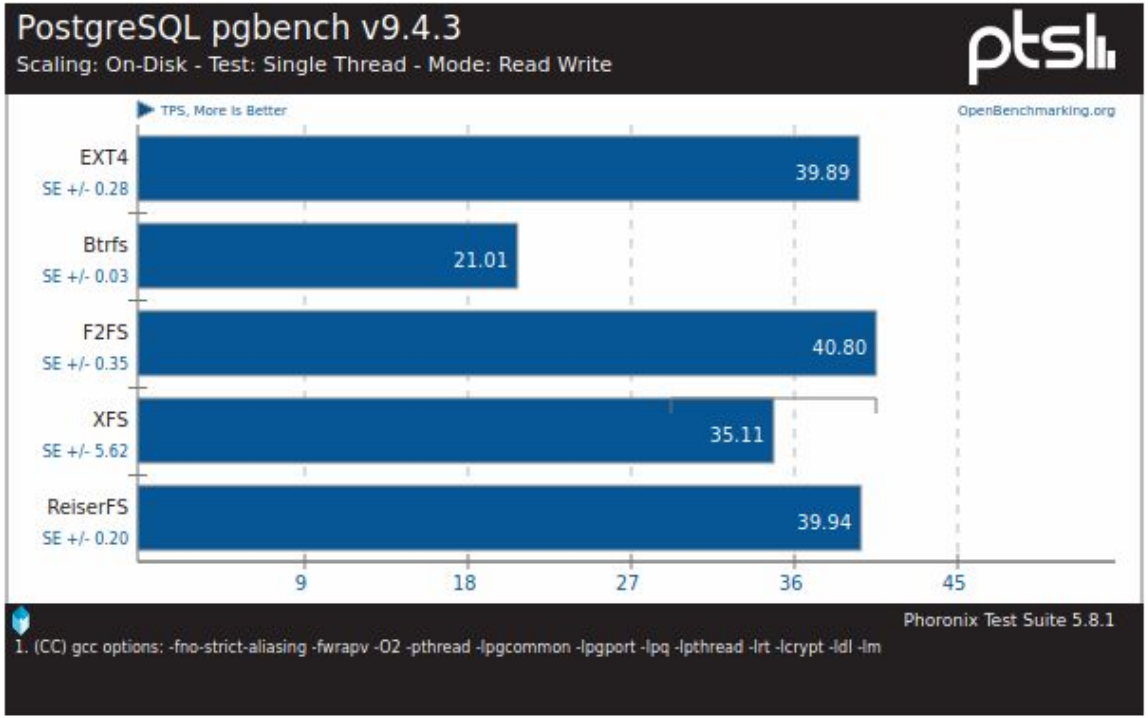

“This is a simple benchmark of PostgreSQL using pgbench.”

For PostgreSQL, F2FS and ReiserFS along with EXT4 tended to be the fastest file-systems while Btrfs was the slowest (note those using Btrfs with databases probably want to set the "nodatacow" mount option).

How Many Megabits per second in a Gigabit per second?

How Many Megabits per second in a Gigabit per second?

1 Megabit/sec is equal to (Gigabit/sec)/1000.

1 Gigabit/s = 1000 × Megabits/sec.

Gbps : Gigabit per second (Gbit/s or Gb/s)
Mbps : Megabit per second (Mbit/s or Mb/s)
prefix mega : 1000^2
prefix giga : 1000^3
1 megabit = 1000^2 bits
1 gigabit = 1000^3 bits
1 gigabit = 1000^{3-2} megabits
1 gigabit = 1000 megabits
1 gigabit/second = 1000 megabits/second
1 Gbps = 1000 Mbps

Source: gmb.org, Gbps to Mbps Converter
https://www.gbmb.org/gbps-to-mbps

Storage Media

1. USB3 is fast
2. SSD is fast too
3. USB2 is really slow
4. Moving Terabytes of Data takes time

Gbps to MB/s Converter

How Many Megabytes per second in a Gigabit per second?

1 Megabyte/sec is equal to (8 × Gigabit/sec)/1000.

1 Gigabit/s = 125 Megabytes/sec.

```
Gbps : Gigabit per second (Gbit/s or Gb/s)
MB/s : Megabyte per second
1 byte = 8 bits
1 bit  = (1/8) bytes
1 bit  = 0.125 bytes
1 megabyte = 1000^2 bytes
1 gigabit  = 1000^3 bits
1 gigabit  = (1000 / 8) megabytes
1 gigabit  = 125 megabytes
1 gigabit/second = 125 megabytes/second
1 Gbps = 125 MB/s
```

Gigabits per second to Megabytes per second Examples

- 10 Gigabit Ethernet speed 10 Gbit/s = 1250 Megabytes per second.
- USB 3.0 transmission speed 5 Gbit/s = 625 Megabytes per second.

Source: Gbps to MB/s Converter
https://www.gbmb.org/gbps-to-mbs

Hard Drive Choices / Prices

The Cost of Storing Data

$119.99

Price valid through 7/3/20 & Free Shipping

Seagate Backup Plus Hub 8TB Desktop Hard Drive with Rescue Data Recovery…

(769)

$64.99

After $5 OFF

Seagate Backup Plus Ultra Touch 2TB Portable Hard Drive with Rescue Data…

(509)

$119.99

Free Shipping

Seagate Backup Plus 5TB Portable Hard Drive with Rescue Data Recovery…

(818)

Source: https://www.costco.com/external-hard-drives.html

Is SMR a scandal or a misunderstood technology?

How does Shingled Magnetic Recording work?

Introducing Seagate SMR

Seagate Shingled Magnetic Recording, or SMR, is breaking barriers to new areal densities and adding greater than 25% capacity growth by maximizing the number of tracks per inch on a single disk.

Traditionally, track spacing shrank with the size of the recording head's reader and writer elements (Figure 1).

Figure 1. Conventional Track Spacing

The reader and writer elements of today's perpendicular magnetic recording HDDs have reached a physical limitation. Without future recording technologies, they cannot become smaller, nor can the tracks they read and write.

SMR achieves higher areal densities by squeezing tracks closer together. Tracks overlap one another, like shingles on a roof, allowing more data to be written to the same space. As new data is written, the drive tracks are trimmed, or shingled. Because the reader element on the drive head is smaller than the writer, all data can still be read off the trimmed track without compromise to data integrity or reliability. In addition, traditional reader and writer elements can be used for SMR. This does not require significant new production capital to be used in a product, and will enable SMR-enabled HDDs to help keep costs low.

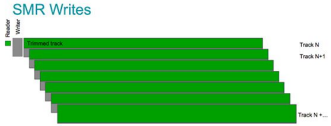

Figure 2. Track Spacing Enabled by SMR Technology

When a user needs to rewrite or update existing information, SMR drives will need to correct not only the requested data, but any data on the following tracks. Since the writer is wider than the trimmed track, all data in surrounding tracks are essentially *picked up* and as a result will need to be rewritten at a later time (Figure 3). When the data in the following track is rewritten, the SMR drive would need to correct the data in the subsequent track, repeating the process accordingly until the end of the drive.

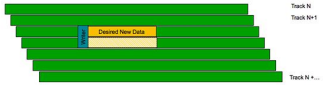

Figure 3. Writer Overlap on Trimmed Tracks

For this reason, SMR groups tracks into bands, where the shingling process stops (Figure 4). This enables an SMR drive to better manage these rewrites. This also improves the drive's write performance by grouping tracks into bands that optimize the number of tracks that need to be rewritten.

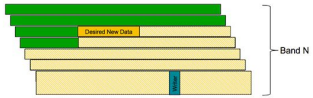

Figure 4. SMR Band Structure

Source: Introducing Seagate SMR
https://www.seagate.com/tech-insights/breaking-areal-density-barriers-with-seagate-smr-master-ti/

Seagate Shingled Magnetic Recording, or SMR, is breaking barriers to new areal densities and adding greater than 25% capacity growth by maximizing the number of tracks per inch on a single disk.
https://www.youtube.com/watch?v=3UFUfv9n420

SMR Magic

- Write heads are bigger than read heads
- Existing technology read/write heads used
- Overlapping tracks on write allow integrity on reading
- However, rewriting data means lots of data relocation

Source:
https://www.seagate.com/tech-insights/breaking-areal-density-barriers-with-seagate-smr-master-ti/

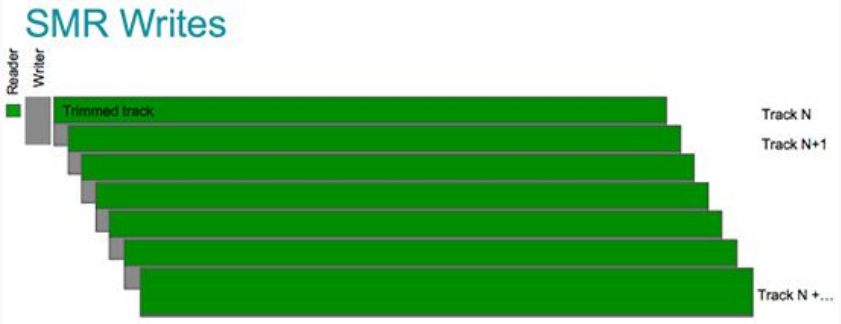

"improves the drive's write performance by grouping tracks into bands that optimize the number of tracks that need to be rewritten."

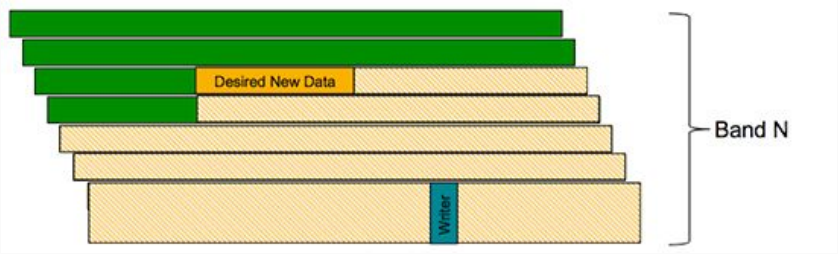

Seagate

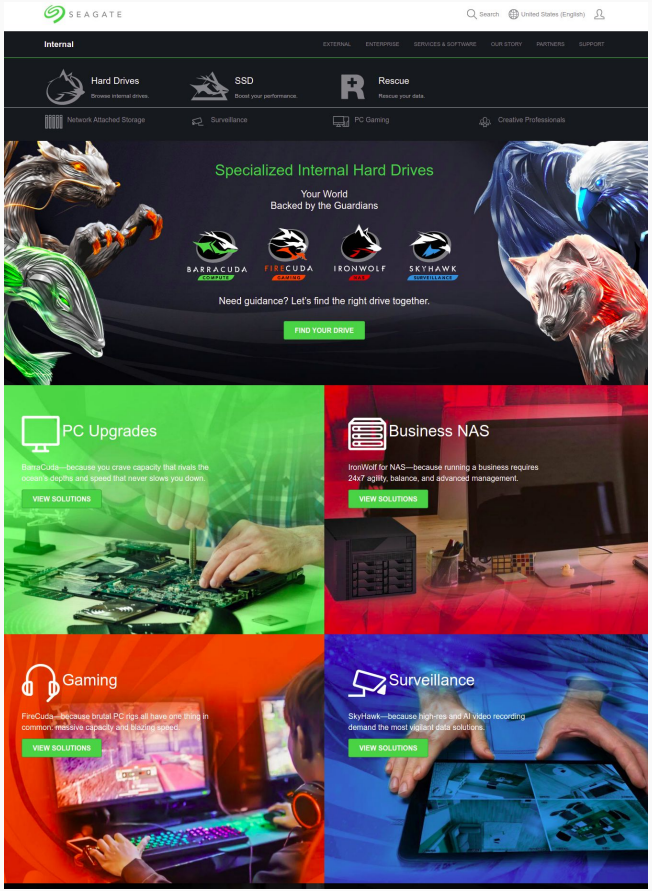

https://www.seagate.com/internal-hard-drives/right-drive/

https://www.seagate.com/internal-hard-drives/

Western Digital

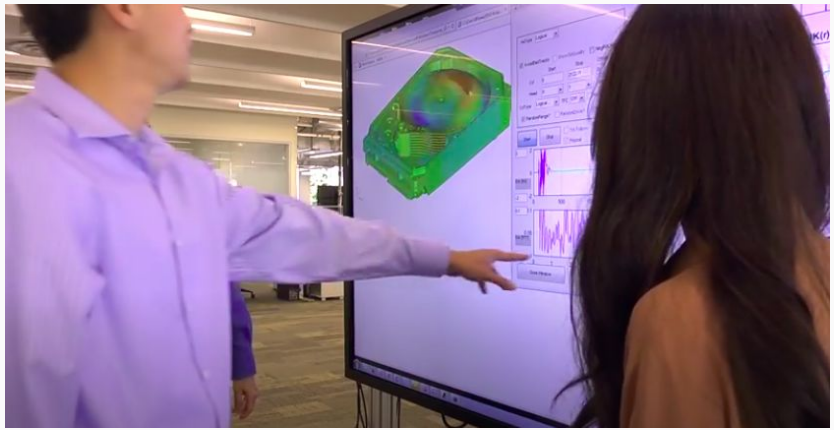

https://youtu.be/9C54Oo8-o1E

https://www.westerndigital.com/products/data-center-drives#solid-state-ssd