

Benchmarking Hard Disk Duplication Performance in Forensic Applications

By Robert Botchek, Tableau, LLC

Introduction

Computer forensics, the practice of capturing and analyzing digital data from many sources, often involves the capture and analysis of data from computer mass storage devices like hard disks. In many forensic practices, the number of disk acquisitions and the amount of data acquired become limiting factors in the throughput and effectiveness of the forensic practitioner.

Many factors influence the speed at which a hard disk transfers data. Moreover, forensic best practices place additional requirements – like disk hashing – on the data acquisition process.

This paper begins with an overview of hard disk capacity and performance trends for the past three decades, using this information to demonstrate the magnitude of the growing challenges facing computer forensic practitioners. After a brief survey of modern disk acquisition methodologies, we continue with a detailed, technical discussion of the factors which dictate "imaging" (or disk duplication) performance, especially as it relates to the requirements of forensic process. Finally, we present real-world data illustrating the performance of typical, widely used forensic tools.

Thus informed, we hope forensic practitioners will have a deeper understanding of the technical factors which directly influence the rate at which they perform disk imaging tasks which are central to their jobs. With this information in hand, forensic practitioners should have well-grounded, realistic expectations for their tools and should be able to make informed methodology and equipment choices in order to optimize their daily work.

Background

The process of making an "image" – or duplicate copy – of a hard disk is central to the practice of computer forensics. Given the vital need to maintain the integrity of the original evidence, forensic analysis is rarely performed on "original" hard disk devices. When analysis is performed on the original devices, it is often more in the context of a fast "triage" process by which the examiner decides whether or not to capture a complete image of the subject device for further analysis.

So, practitioners of computer forensics often find themselves imaging hard disks and similar, large computer storage devices. The quantity of data involved is staggering. A "small" modern hard disk may have 40 *billion* bytes of data, with more typical devices having 80 – 300 billion bytes of data. As of this writing, the largest individual hard disks have capacities of 1 *trillion* bytes, and capacities will easily continue to grow 20 – 50% (or more) per year. Individual cases may involve numerous hard disks, further exacerbating the data collection challenge.

Figure 1 illustrates one factor underlying rapid growth in storage capacities, growth in hard disk capacities vs. cost (referred to as "Megs per \$" in the figure and today more commonly referred to as the cost per gigabyte)¹. As seen in the figure, capacity has a sustained history of rapid growth compared with cost, and this encourages users to purchase and use larger amounts of storage with each passing year. Later in this paper, Figure 2 illustrates that capacity, independent of cost, is growing at an exponential rate.

¹ Source: <http://www.mattscomputertrends.com/harddrives.html>.

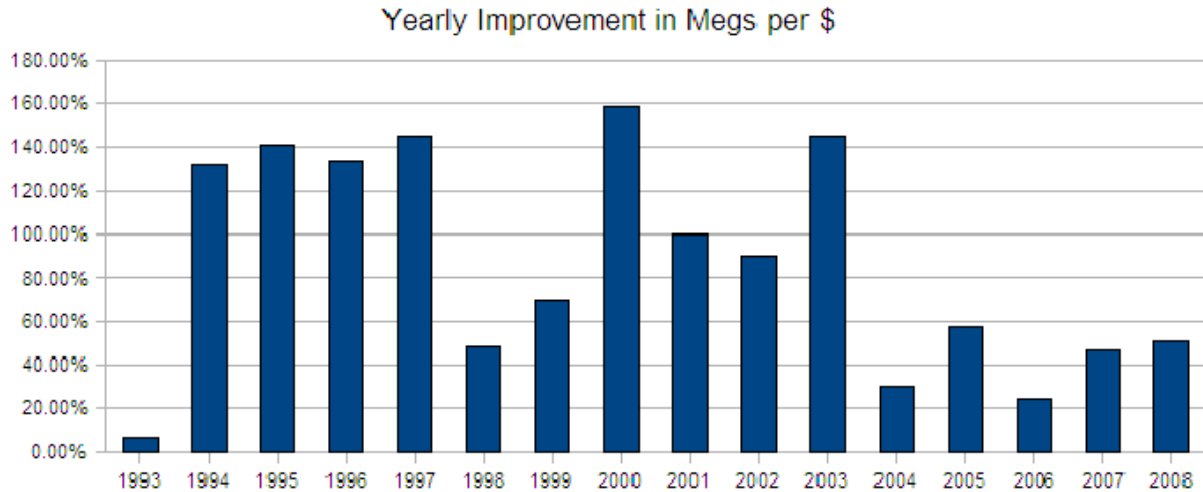


Figure 1. Growth in Hard Disk Capacity vs. Cost

Given the enormous quantities of data involved, forensic practitioners are constantly seeking faster and more reliable ways to duplicate evidence devices. Even more important than speed, however, forensic practitioners require the ability to authenticate and re-authenticate the data which is captured for analysis. Forensic practitioners commonly generate a "digital fingerprint" for the data captured from each hard disk, and the digital fingerprint is most often generated by passing each hard disk image through a "hash" calculation function.

Hashes are computer/mathematical algorithms such as "MD5" and "SHA-1" – often associated with cryptography – which generate seemingly random fingerprints (called "message digests") based on a set of source data. Given the same source data, a given hash function will always generate the same fingerprint. The unique property of the hash functions used in forensics is that it is "impossible" to construct a second data set which will produce the same hash fingerprint as an original data set². In forensics this is relevant because it provides a way for practitioners to re-verify the data periodically and to demonstrate that the data remains a faithful copy of the original evidence and that no tampering, intentional or otherwise, has altered the data.

Thus, when making copies of original evidence, forensic practitioners need computer hardware and software which can move enormous quantities of data and which can compute hash fingerprints for that data as quickly as possible.

² Researchers have found ways to generate two or more data sets which produce identical hash message digests, called hash "collisions". However, these attacks rely on the ability to control each original data set. At present, there are no known, computationally feasible methods for generating a second data set for which the hash fingerprint matches an original data set not generated by the researcher.

Equipment and Procedures

Forensic practitioners typically use several methods to acquire – or image – a hard disk device.

Write blocker based acquisition. The subject device is connected to a hardware "write blocker" and the write blocker is connected to a host computer. The write blocker allows the host computer to read a faithful image of the data on the hard disk while preventing the host computer from modifying the contents of the hard disk. Software on the host computer is responsible for ordering the reading of data from the hard disk through the write blocker. That same software is generally capable of calculating one or more hash fingerprints while copying the data. The copy itself can go to any of numerous destinations, including another hard disk attached to the computer or even a network storage device.

Write blockers are typically the most versatile method of connecting to and imaging a subject storage device. Hardware write blockers are available today for IDE, SATA, SCSI, and USB device interfaces³, offering rich options for attaching subject devices of many types to the host computer used for acquisition and analysis. On the "host side" of the write blocker, FireWire and USB interfaces offer robust, mature "hot-swappable interfaces" to the host computer. This means that many different evidence devices can be connected to a host computer throughout the day without requiring time-consuming re-boots of the host computer. This also means that the full range of software on the host computer is available to preview or image a given hard disk through the write blocker.

The latest generation of write blockers are beginning to offer eSATA as another "host interface" option. eSATA has the promise of faster transfer rates than FireWire and USB. Unfortunately, the host device drivers for eSATA often do not have the mature, robust support for hot-swapping offered by FireWire and USB host interfaces. The lack of robust host-side eSATA driver support will tend to delay the wide-scale adoption of eSATA as a host-side interface for write blockers for perhaps one to three years.

When using a hardware write blocker numerous factors can limit overall imaging performance, including the internal design of the host computer (memory controller, internal PCI bus design, etc.) and the performance of the device/interface to which the copy of the data is being written. These kinds of factors greatly complicate the analysis of write blocker based forensic imaging. For this reason, we do not study write blocker based imaging further in this paper. Even so, many of the observations about disk performance later in this paper are equally applicable to write blocker based processes.

Duplicator based acquisition. The subject device is connected to a dedicated hardware duplicator. The duplicator is, in essence, a small specialized computer which generally combines the functions of the write blocker and host computer as described in the previous paragraph. As such, the duplicator orders the reading of data from the source hard disk, performs hash calculations on the data, and writes a copy (or sometimes multiple simultaneous copies) of the data to other hard disks. Duplicators may be faster than write blockers, but only if they have a fast internal "data path" and the ability to calculate hash fingerprints very quickly. Duplicators tend to be more limited in their options than PC-based acquisitions. For example, an examiner using a write blocker connected to a host computer can "preview" a hard disk before copying it, using the rich user interface and software capabilities found on modern computers.

Software based acquisition. Software based acquisition is similar to write blocker based acquisition, except that the subject hard disk is connected directly to the host computer which will read it. Modern operating systems will often write automatically to a directly attached hard disk, so software based acquisition relies on custom software which *will not* write directly to the subject hard disk. For example, there are numerous Linux-based "boot CDs" which can be used to boot a subject computer and examine / copy the hard disks on that computer.

³ Hardware write blockers for other interfaces are expected in 2009 and beyond.

Software based acquisition lacks some of the key advantages of both write blocker and duplicator based processes. Most notably, the host computer must have the correct device interface for each device to be imaged and much be rebooted each time the device is changed. This can be particularly troublesome when imaging more esoteric devices such as SCSI hard disks. Further, the most popular forensic analysis tools run under Microsoft Windows and not under the Linux platforms on which most software based acquisition platforms are built.

When considering both the advantages and limitations of software based acquisition, we observe that software based acquisition is often employed in situations where the physical removal of the hard disk from its host is difficult or impractical given time constraints. The hard disks in notebook computers often fall into this category. Time constraints may be apparent in operations which involve "screening", such as operations at immigration and customs checkpoints.

In practice, most forensic practitioners will use a combination of several equipment and procedural choices, picking the equipment and procedure most appropriate for the forensic examination at hand.

In the remaining sections of this paper we will discuss factors which affect hard disk performance and which consequently affect imaging performance. These conclusions have ramifications across equipment and methodology choices, though we will tend to focus on duplicator based imaging in this paper.

Hard Disk Capacity and Performance Trends

Just as hard disk capacities vary widely and grow quickly over time, hard disk performance also varies widely, though performance has grown more slowly over time than capacity. The following figure illustrates the growth in hard disk capacities for the past three decades⁴.




Figure 2. Hard Disk Capacities (1980 - present)

In the mid 1980's we can see that hard disks had typical capacities of 5MB – 30MB (0.005 – 0.03GB). At that same time, typical disk media rates were measured in hundreds of kilobytes per second. Today, with typical disk capacities of 100GB – 1TB, typical disk media rates are measured in the range of 30MB/sec – 130MB/sec.

So, consider the comparison in capacity vs. performance. Over the 25 year period from the early/mid 1980's to the present, consumer hard disk capacities have increased by approximately the ratio of 1TB / 30MB or roughly a factor of 30,000+. Over that same period of time, disk media rates have increased by roughly the ratio of 100MB/sec / 100kB/sec or a factor of only 1,000. This has one very obvious consequence: The theoretical time required to image a complete hard disk has grown approximately 30-fold in the past 25 years.

⁴ Source: http://commons.wikimedia.org/wiki/Image:Hard_drive_capacity_over_time.png.

Understanding Hard Disk Performance in Forensic Applications

The preceding discussion of hard disk capacities and media transfer rates grossly oversimplifies the discussion of hard disk imaging time. In order to gain a better understanding of hard disk imaging time we need to begin with some common terminology and technical concepts.

In the preceding section we used the term "media transfer rate". "Media transfer rate" is the term we will use herein to discuss the rate at which the hard disk actually reads or writes data to the physical media. In the case of solid state drives which are now coming to market, the media transfer rate would be the rate at which the solid state drive electronics can move data to/from the non-volatile memory (generally FLASH).

There are many other metrics used to describe hard disk performance, including *interface transfer rate*, cache buffer sizes, average seek time, track-to-track seek time and rotational latency. The concepts of interface transfer rate and cache size are also applicable to solid state drives, though discussion of seek time and rotational latency is obviously not relevant to solid state drives.

In the forensic process of imaging a subject hard disk we are generally reading the *entire* hard disk from the beginning to the end, and we are doing so sequentially. Since we are reading the hard disk sequentially, the average seek performance of the hard disk is not relevant...the imaging process simply doesn't request random seeking across the source media. Track-to-track seek performance *is* relevant as we must sequentially read and write all sectors, and that requires the source and destination drives to step linearly from one track to the next. However, track-to-track seek time is often on the order of 1 millisecond or less, and does not introduce a large degradation in typical imaging performance.

The sequential nature of forensic imaging also largely eliminates the benefits of fast *interface* transfer rates and large on-disk cache buffers. Even experienced forensic practitioners are often fooled by claims of high *interface* transfer rates, so we will spend some time discussing the difference between media transfer rate and interface transfer rate.

Interface	Peak Data Rate (bits/sec)	Peak Data Rate (bytes/sec) ⁵
IDE UDMA 5	-	100 MBytes/sec
IDE UDMA 6	-	133 MBytes/sec
SATA I	1.5 Gbits/sec	150 MBytes/sec
SATA II	3.0 Gbits/sec	300 MBytes/sec
eSATA	Same as SATA I or II	Same as SATA I or II
Ultra SCSI 320	-	320 MBytes/sec
FireWire400 (1394A)	400 Mbits/sec	40 MBytes/sec
FireWire800 (1394B)	800 Mbits/sec	80 MBytes/sec
USB 2.0	480 Mbits /sec	48 MBytes/sec

Table 1. Peak Interface Transfer Rates for Selected Interface Technologies

Common hard disks today employ the SATA (Serial ATA) interface standard. As shown in Table 1, SATA I offers an interface transfer rate of 1.5Gbit/sec and SATA II offers an interface transfer rate of 3.0Gbit/sec. To those not familiar with the difference between media transfer rate and interface transfer rate, these numbers make it seem like a SATA I drive should be able to transfer data at 150MB/sec or a SATA II drive should be able to transfer data at 300MB/sec.

⁵ Many serial interface technologies (SATA, FireWire800) use "8b/10b" encoding. So, one byte is equal to 10 bits transferred across the serial interface.

These peak interface transfer rates *are* relevant when the disk device can take advantage of its built-in cache to buffer the data. For example, when processing a read request from the host, most hard disks will internally read more data than actually requested by the host, buffering the balance of the data in the built-in cache. Subsequently, if the host requests the next sequential blocks of data from the disk, the drive already has the data buffered (or partially buffered) in cache and can transfer the data at the peak rate supported by the disk interface. Similarly, when a drive is operating in "cached-write" mode (where the drive will accept the data and complete the write command before actually recording the data on the physical media) the drive can accept write data from the host computer at the peak interface transfer rate and can buffer that data in cache.

However, in sequential forensic imaging of a drive, it is not likely that the caching behavior of the drive would enable the drive to use the peak interface transfer rate. If the host computer or duplicator driving the duplication process is able to transfer data faster than the *media* transfer rate of the drive, then the source drive's cache will always be empty on back-to-back sequential reads and the destination drive's cache will always be full on back-to-back sequential writes. Considering the case of reads from the source, it doesn't matter how fast the peak interface transfer rate is if the drive can only provide the data at the media transfer rate. Likewise in the case of writes to the destination, once the drive's cache is full, the drive can only accept additional write data as fast as it can flush earlier write data to the media, again limited by the media transfer rate.

Conclusion: Media transfer rate is generally the rate-limiting factor in determining how fast a hard disk can source or accept data during the long, sequential operations typical of forensic drive imaging. There are exceptions to this rule, as we will discuss in a later section, but as a general statement, the media transfer rate of a hard disk or solid state drive is the primary drive-controlled factor to consider when studying imaging performance.

Media Transfer Rates in More Detail

In the preceding section we establish the basis for claiming that media transfer rate is the primary drive-controlled factor in limiting disk imaging performance. So, what determines the media transfer rate of a drive?

First, let us consider the more simple case of solid state drives. A solid state drive, unlike a traditional hard disk, has no moving parts. Where a hard disk has one or more rotating platters on which the data is recorded magnetically, a solid state drive is typically based on non-volatile memory such as FLASH. In the case of a solid state drive the media transfer rate is closely related to the rate at which the non-volatile memory can be read or written. Non-volatile read and write data rates are often quite different, with the write rates being significantly slower in many cases than the read rates. So, a solid state drive is likely to have two different media transfer rates, one for reading and one for writing, and those rates are likely to be constant across the entire addressable range of the device.

Rotating media hard disks present many more variables which affect the media transfer rate. Hard disk media transfer rates are principally affected by:

The rotational rate of the drive. As the drive rotates more quickly, the data bits on the magnetic platters pass under the read/write heads more quickly, increasing the media transfer rate. 7,200RPM drives spin 33% faster than 5,400RPM drives. Thus, for a given bit-density (also called recording technology), a 7,200RPM drive should have a media transfer rate 33% higher than its 5,400RPM counterpart. A 10,000RPM or 15,000RPM drive will be proportionally faster, while a 4,200RPM notebook hard disk will have among the slowest media transfer rates.

The bit-density or recording technology. Most of the gains in hard disk capacity in the past two decades have come from increases in "aerial bit density", or the number of bits that can be recorded over the surface area of each hard disk platter. Aerial bit density can increase either by reducing the spacing of bits along each track or by decreasing the spacing between tracks. Interestingly, those increases in aerial bit density which result from reduced bit spacing along a track also drive increases in media transfer rate. Consider: If two models of hard disks are rotating at the same rate, but one model packs twice as many bits along each track compared with the other model, then twice as many bits will pass under the read/write head in the higher-capacity drive in the same amount of time.

Advances in head and "read-channel" technology. The read/write head and the electronics which control it will also limit the media transfer rate. Gains in aerial density and rotational rate will not improve overall media transfer rates unless the read/write head and control electronics are able to keep up.

Reading vs. writing. For reasons that are beyond the scope of this paper, some hard disks read and write data at very different rates, with reads often being the faster of the two. For example, a modern Samsung 1TB hard disk typically reads data about 20% faster than it can accept data to be written to the disk.

The position of the data on the drive. Media rates are *not* constant over the surface area of a rotating media hard disk. The "beginning" of the hard disk (i.e., sector 0) is typically located on the outermost user-addressable track on the hard disk, while the "end" of the hard disk (the highest addressable sector) is located on the innermost user-addressable track (see Figure 3). Hard disks, unlike CD-ROMs, rotate with a constant angular velocity. So, whether the read/write heads are positioned over the outermost or innermost tracks (illustrated in red in Figure 3), the angular rotational rate remains the same, for example, 7,200RPM.




Figure 3. Inner and Outer Track Radii

In old hard disk designs (prior to the early 1990's), a "sector" occupied the same angular distance on a track, regardless of the position of the track on the platter. That meant that the bits on inner tracks were spaced more closely than bits on outer tracks. In the late 1980's and early 1990's hard disk designers realized that a lot of capacity was being wasted on the outer tracks because of the increased spacing between bits. So, with the advent of "zoned-bit encoding", disk manufacturers started taking advantage of the longer "outer" tracks by encoding more sectors onto those tracks. In an optimal, modern hard disk, bits are always recorded at approximately the same spacing on the media. So, there are progressively more and more sectors on each track as one moves from the inside of the platter to the outside.

What most computer users don't realize is the profound impact this has on media transfer rates. We can see that the outermost track is longer than the innermost track, and therefore has more sectors. How many more? Roughly speaking, the number of sectors on the outermost track is greater than the number of sectors on the innermost track in proportion to the radius of each track as shown below.

$$\begin{aligned} \text{Length of innermost track, } D_1 &= 2\pi R_1 \\ \text{Length of outermost track, } D_2 &= 2\pi R_2 \end{aligned}$$

If D_S is the average length of a sector on a track, then.

$$\begin{aligned} \text{Number of sectors on innermost track, } S_1 &= D_1 / D_S = 2\pi R_1 / D_S \\ \text{Number of sectors on outermost track, } S_2 &= D_2 / D_S = 2\pi R_2 / D_S \end{aligned}$$

And the ratio of the number of sectors on the outermost track vs. the innermost track is:

$$S_2 / S_1 = (2\pi R_2 / D_S) / (2\pi R_1 / D_S) = R_2 / R_1$$

In Figure 3 we can see that R_2 is approximately twice R_1 , leading one to conclude – correctly – that a modern hard disk with platters like that of Figure 3 will have approximately twice as many sectors on the outermost track than on the innermost track.

Remembering that the disk is rotating at a constant rate, we can conclude that the ratio of $S_2/S_1 = R_2/R_1$ also indicates how much *faster* the media transfer rate will be at the outermost track when compared with the innermost track. Again, in this illustration, the ratio is approximately 2:1.

Figure 4 and Figure 5 are screen shots from the HD Tune utility running under Microsoft Windows. HD Tune has the ability to graph the effective transfer rate across the entire surface of the hard disk. The difference in media transfer rates across the surface of a hard disk is readily apparent for each of the two drives (blue curve in each figure).

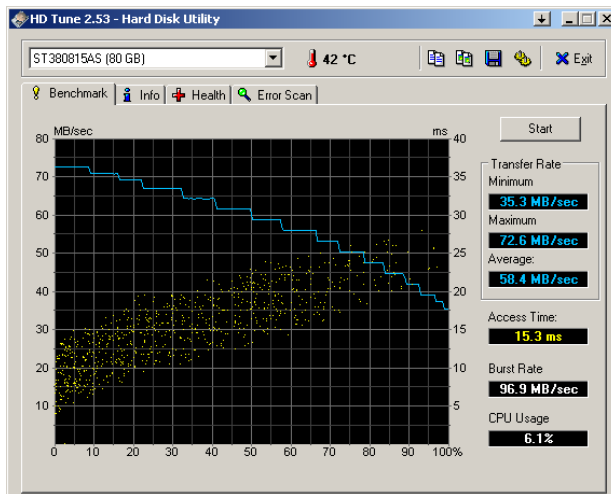


Figure 4. Hard Disk Performance - Seagate ST380815AS (Barracuda 80GB SATA)




Figure 5. Hard Disk Performance - WD WDC3000GLFS (VelociRaptor 300GB SATA)

In the case of the Seagate 80GB Barracuda drive, the ratio of performance at the beginning of the drive is roughly 2.08 : 1. The corresponding ratio for the Western Digital 300GB VelociRaptor is roughly 1.54 : 1. In other words, for a representative sample of modern drives, the media transfer rate at the beginning of the drive is 50 – 100% greater than that at the end of the drive.

The preceding figures also verify a conclusion presented in an earlier section, namely, when transferring data sequentially across the entire drive – as during forensic imaging – the most important drive-controlled factor is the media transfer rate. Both of the hard disks shown in the HD Tune results use SATA interfaces, and therefore have peak interface rates of at least 1.5Gbit/sec. At no point does the transfer rate of either hard disk approach the 150MB/sec theoretical performance of the SATA interface.

Conclusion: It is impossible to image an entire hard disk at a rate faster than the drive's media transfer rate (specifically its media transfer rate for reading). It is impossible to sequentially transfer to a hard disk faster than the drive's media transfer rate for writing (which may be different from its read rate). The computer, write blocker, duplicator and software can be *infinitely fast* and the answer will still be the same. Both in theory and in practice, performance cannot exceed the media transfer rates for the hard disk being imaged and the storage device to which the image is being written.

When Media Transfer Rate Does Not Dictate Forensic Imaging Performance

Having explained why forensic imaging can never exceed the media transfer rate of the source and destination hard disks we should also explore situations in which the effective imaging rate is *slower* than the media transfer rate. Stated simply, the process of creating a forensic image – copying the contents of a subject device to a destination and calculating one or more hash fingerprints – can only be as fast as the slowest step in the process.

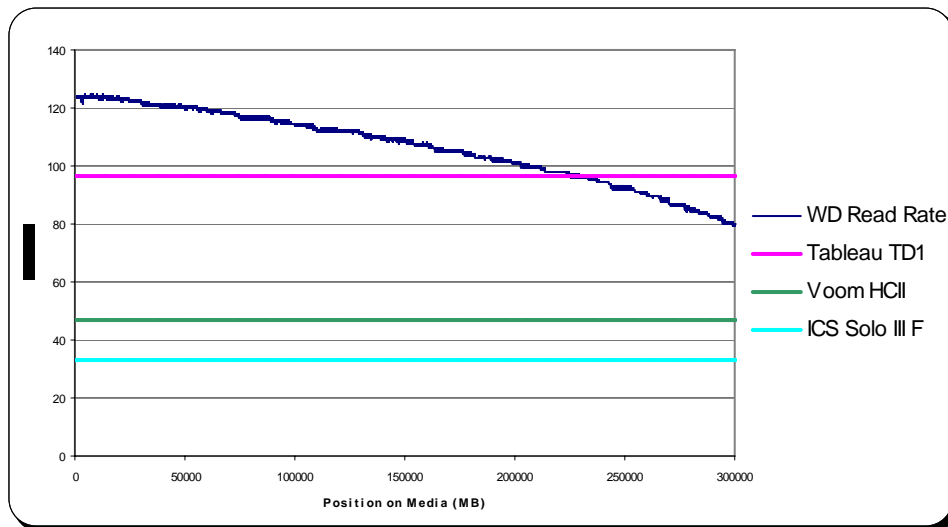


Figure 6. WD VelociRaptor Media Transfer Rate vs. Duplicator Performance

Figure 6 presents the media transfer rate of the Western Digital 300GB VelociRaptor (blue line⁶) and compares it with the measured performance of several leading forensic duplicators, the Tableau TD1 (magenta line), Voom HardCopy II (green line), and ICS Solo III F (light blue line). Each duplicator's performance is shown as a straight line, indicating that the duplicator's performance does not change as the imaging process traverses the hard disk surface. Instead, duplicator performance depends on the internal design of the duplicator's data path and on the design of its hash calculation circuitry (or firmware).

⁶ The data of Figure 6 is nearly identical to that of Figure 5. Data for Figure 6 was collected using the HD Speed benchmarking utility.

The performance levels shown for each duplicator are based on real-world measurements under controlled settings. Each duplicator was used to image the same hard disks, and timing/performance measurements were made independently of the values displayed by the duplicators themselves. Of key significance, each duplicator performed the duplication with hash calculations enabled. While the duplicators are typically capable of copying data without performing a hash calculation⁷ – and will generally achieve higher performance when hashing is disabled – a forensic practitioner will rarely use the duplicators without hashing enabled given the requirement to generate a hash fingerprint for each forensic image.

So, looking at the data in Figure 6 we can draw several conclusions. When duplicating a very high-performance disk like the WD 300GB VelociRaptor, the slower duplicators (HCII and Solo III F) will be limited by their internal designs and will not approach the media transfer rate of the VelociRaptor. For the slower duplicators, duplication time can be accurately estimated by dividing the 300GB capacity by the transfer rates of each duplicator (roughly 2.0GB/sec and 2.8GB/sec, respectively), shown again as straight lines in Figure 7.

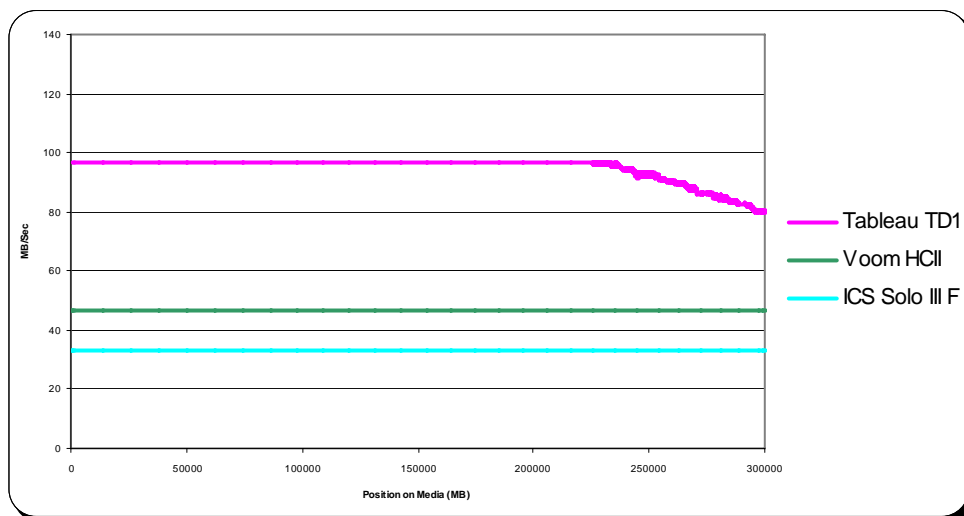


Figure 7. Effective Duplication Performance with WD 300GB VelociRaptor

The combined performance of the VelociRaptor and the Tableau TD1 is more complex. For roughly the first 225GB of the VelociRaptor's 300GB capacity, the media transfer rate of the VelociRaptor is higher than the sustained 96.5MB/sec capability (data path + hash calculation) of the TD1 (straight magenta line in Figure 7). However, roughly 225GB into the drive, the VelociRaptor's media transfer rate falls below the transfer rate of the TD1. At this point the effective imaging performance is limited by the VelociRaptor (sloped magenta curve in Figure 7). So, for the first 225GB the TD1 averages roughly 5.8GB/min and for the last 75GB the TD1 is limited by the VelociRaptor and averages roughly 5.1GB/min.

As of this writing the WD 300GB VelociRaptor is one of the fastest hard disks in the world. So, we also present results for a hard disk which is more representative of drives commonly being acquired in forensic examinations today, the Seagate 80GB Barracuda.

⁷ Hashing is always enabled in the Tableau TD1. The TD1 uses a patent pending hashing system which operates at the same speed as the TD1's internal data path, so hashing does not degrade TD1 performance.

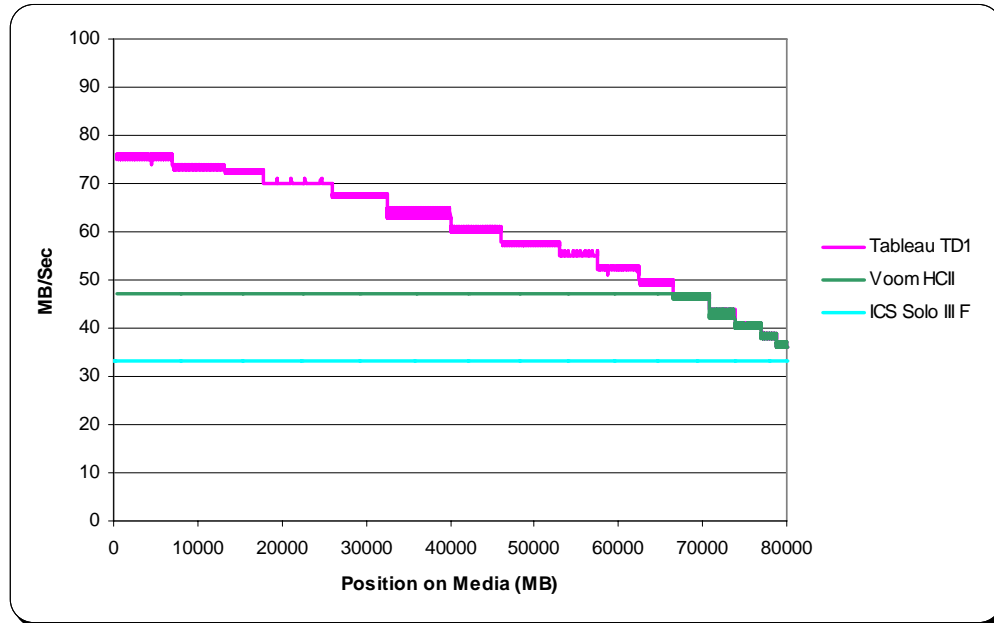


Figure 8. Effective Duplication Performance with Seagate 80GB Barracuda

Once again, the performance of the slowest duplicator, the Solo III F, is dictated purely by limitations of its internal design. Likewise, for most of the duplication the HCI is also limited by its internal design, though roughly 65GB into the drive it, too, becomes drive limited. The Tableau TD1 is limited by the Seagate's media transfer rate throughout the entire imaging process.

This is an important point which is commonly misunderstood by forensic practitioners: Even very fast duplicators like the TD1 can go no faster than the drives they are imaging.

Conclusions

Tableau developed this paper in order to educate members of the forensic community and to correct many commonly held misunderstandings regarding the performance of mass storage devices and the forensic tools used to image those devices.

We have illustrated many of the real-world factors which affect the rate at which forensic tools are capable of creating forensically sound images of subject devices. An understanding of these factors should help forensic practitioners to set realistic expectations for their tools, to configure their operating environments and their tools effectively, and to determine if their tools are operating optimally.

About the Author

Robert Botchek is the President and founder of Tableau, LLC. Tableau is a world leader in forensic products including hardware write blockers, duplicators, and hardware accelerators.

Mr. Botchek began his professional career in 1984 at Trantor Systems, Ltd. where he began working with the predecessors to many of today's modern storage interface technologies. As VP Engineering at Trantor, Mr. Botchek led the development of numerous products for I/O technologies related to hard disks, hard disk RAID's, CD-ROMs and tape drives. After selling Trantor Systems to Adaptec in the early 1990's, Mr. Botchek became the VP of Business Development for Electric Gravity; and after selling Electric Gravity to Microsoft in the mid 1990's, Mr. Botchek returned to Adaptec to hold senior executive positions in 1394 (FireWire) and FibreChannel groups.