

KARELIA UNIVERSITY OF APPLIED SCIENCES
Degree Program in Information and Communication Technology

Pekka Kinnunen

TEST ENTITY FOR MICROSOFT RESILIENT FILESYSTEM

Thesis
October 2019



THESIS
October 2019
**Degree Program in Information and
Communication Technology**

Karjalankatu 3
80200 JOENSUU
FINLAND
+ 358 13 260 600

Author (s)
Pekka Kinnunen

Title
Test Entity for Microsoft Resilient Filesystem

Commissioned by
Blanco Technology Group IP Oy

Abstract

The purpose of this thesis was to test Microsoft's latest filesystem called Resilient Filesystem (ReFS) that was originally developed to replace current filesystem that is widely in use – NTFS. Since the new filesystem is not used by as many people as NTFS, there is a lot less information about it. The focus is on testing if there are any issues that are caused by the wide variety of features this new filesystem has.

Theory portion focuses on the importance of secure data erasure, in other words, why is it important to make sure data is completely gone from hard drives, what does it entail and what kind of standards there are in data sanitization. It will also mention how EU GDPR relates to data erasure

Results on the defined scenarios showed that some caution should be exercised when dealing with ReFS. There are some features that makes data erasure a bit more difficult, but not impossible. This thesis resulted in meaningful data for the stakeholders and great experience for the person conducting the tests.

Language

English

Pages 51

Appendices 1

Pages of Appendices 1

Keywords

secure data erasure, data sanitization, ReFS, data recovery

Contents

1	Introduction	5
2	Importance of secure data erasure	6
2.1	Effects of data collection	6
2.2	GDPR of EU	7
2.3	Data in hard drive	8
2.4	Secure data erasure	9
2.5	Standards	10
3	Preparing for the assignment	11
3.1	Introduction to Microsoft Resilient Filesystem (ReFS)	11
3.2	Creation of test data	14
3.3	Setting up a test environment for ReFS	17
3.3.1	Test setup requirements (SW/HW)	18
3.3.2	Creating a bootable USB with Windows Server 2016 image	19
3.3.3	Windows Server Setup	20
3.3.4	Installing VMware Workstation	22
3.3.5	Setting up virtual test environment for ReFS	24
3.3.6	Problems encountered	28
4	Test scenarios	30
4.1.1	Scenario 1: Testing File Integrity and erasure	32
4.1.2	Scenario 2: Erasing all the data from ReFS volume	41
4.1.3	Scenario 3: What is the recovered data?	45
4.1.4	Results	47
4.2	Test conclusion - current hypothesis	47
5	Conclusion	49
	References	50

Appendices

Appendix 1 Example report from file erasure software

Abbreviations

GDPR	General Data Protection Regulation, EU's latest data privacy regulation.
ReFS	Resilient FileSystem, Microsoft's new developed filesystem that was supposed to replace NTFS.
IoT	Internet of Things, describes computing devices with ability to transfer data without requiring human interaction.
SCSI	Small Computer System Interface, standard for transferring data between devices.
OS	Operating System, low level software that manages computers resources and provides common services.
ESXi	Elastic Sky X Integrated, type-1 hypervisor. Virtualization platform used to host virtual machines.
NTFS, FAT	New Technology File System, File Allocation Table, different filesystems. NTFS being mostly used in consumer computers.
chkdsk	System tool to check integrity and health of a system volume. Also fixes errors.
USB	Universal Serial Bus, standards for cables, connectors and protocols.
GUI	Graphical User Interface, form of user interface which offers icons and audio.
BitLocker	Full volume encryption provided by Windows operating system. Designed to protect data.
RAID	Redundant array of independent disks. Storage virtualization technology that combines multiple physical disks into logical units.

1 Introduction

In this thesis work, I will go through the assignment that was given to me by Blancco Technology Group IP Oy's research team. This paper will focus on data erasure and recovery technologies that are currently in everyday use. Blancco itself as a company focuses on data erasure.

As I have been working for Blancco, I have become familiar with data erasure, its importance and how it is securely accomplished. As my studies focused more on hardware and systems, so far it has been a good fit for my degree.

In the theory portion, I focus on data erasure in general. What it means to erase data from a hard drive and why it is important. To explain that, I use EU's latest data regulation GDPR as an example and go through some typical misconceptions people might have regarding secure data erasure.

After that I focus on the assignment that was given to me – how does data erasure work on Microsoft latest File System – Resilient Filesystem (aka ReFS). The goal of this thesis work was to find any possible issues this new file system might bring, then possibly to find solutions to the said issues and if this file system should be used in general.

Most of the thesis work's testing and execution was carried out during Spring 2019.

2 Importance of secure data erasure

During the past few years there has been an increasing amount of talk about data collection and storage. Mainly Facebook has been under heavy scrutiny regarding data processing, when it has been revealed where they have used their customer's data.

In addition, internet of things (IoT) has brought more smart devices around us to collect more and more data and to analyse it. These devices can be used to offer new services to our day-to-day life, but they also bring new risks when you consider information security.

Because of these things, the EU has also taken a tougher stance regarding data processing and come up with new privacy regulation, GDPR (General Data Protection Regulation). This thesis focuses on secure data erasure from a computer's hard drive.

2.1 Effects of data collection

Computers, IoT and all kind of data collection have increased the amount of data we are collecting. Every day we come up with new ways to collect, analyse and use the data we have and provide new types of services to ease our lives.

Because there is a lot of data around us, we also must build more data centres to store it. When building data centres, they need to be secure and follow various standards. Especially if customer or sensitive company data is being stored, you need to make sure it does not fall into wrong hands. Depending on the amount of data and what it concerns, in the wrong hands it can be used against the company, sold to the highest bidder or other ways cause irreparable damage to the company.

The above mentioned points bring up the importance of data erasure. When throwing away either broken computers or, for example, changing hard drives from data centre's servers, you need to be sure that no recoverable data is left

on those disks. This can be easily overlooked when throwing away a broken computer thinking since it is broken, no one can use it. When in reality it has always had a functioning hard drive and the faulty part has been somewhere else. Hence, the next person who sees the computer can only detach the hard drive and take a look at the data it contains.

2.2 GDPR of the EU

GDPR comes from the words General Data Protecting Regulation and like the name implies, it focuses on data protection and privacy for all individual citizens within the European Union. It was enforced in May in 2018. Businesses and public bodies have been given two years to prepare for the changes. (Burgess, 2019)

How do EU GDPR and secure data erasure relate to each other? After EU GDPR, consumers have more power over their own data that, for example, service provider has collected. Consumer can ask to view the data or ask it to be completely removed from service provider's databases. Another example could be that a company stops its production in one continent and decides to move all the data they have there to somewhere else.

In these situations, one needs to be sure that there is no company confidential data or consumer data on hard drives when they are disposed of or recycled. Especially, there must not be any personal data relating to an identified or identifiable natural person. This kind of data includes names, home addresses, pictures, email addresses or even IP-addresses. Definition of this in full length is in EU GDPR regulation 2016/679 document.

2.3 Data in hard drive

Theory

A regular mechanical hard drive writes data to a spinning magnetic disk one bit at a time. As a bit can only be 1 or 0, it means the data in the hard drive is saved as one or zero. Hard drive's read/write head reads the data from the surface of the spinning disk and since it is based on magnetism, it retains all the data even without continuous electric current. (Woodford, 2019)

Currently it may be considered to be outdated technology, but it has still retained its popularity mainly because of its bigger capacities and lower price.

Misconceptions

Many casual day-to-day computer users might not think too much of secure data erasure. Especially if they have no technological background, and experience only of day to day use of Windows operating system. All Windows users are probably familiar with its recycle bin, where all user files deleted by the user end up at first. After a while, when the user decides to empty it, he or she can right click its icon and select "Empty Recycle Bin". Doing so will result in this notification:




Figure 1. Emptying Recycle Bin

The number changes depending on the contents of the recycle bin, but the point is the fact that Windows mentions if the user wants to delete the selected files "permanently". That is actually a bit misleading as the files do disappear from the reach of Windows OS, but not permanently from the hard drive.

That is because when you empty the recycle bin, data does not disappear from the hard drive, but it gets marked as rewritable space. This means hard drive is told “whatever was here, you may overwrite”. In other words, data remains in the hard drive until it has been properly overwritten by new data.

Another misconception about data in hard drives is that by formatting the hard drive, you also delete all the data it has. It is true again that the operating system will not be able to see the data it once had after it has been formatted. However, the truth is that in case of formatting a hard drive for operating system to use, it merely removes the address tables. You could compare this to a book. If the hard drive was a book, by formatting it, you are only getting rid of the book’s table of contents page. The data is still there, but finding it now became increasingly harder.

So, until the data gets overwritten, with third party applications, it is still possible to recover it. In many cases in fully usable condition.

2.4 Secure data erasure

When we are talking about secure data erasure or data sanitization, you mean data erasure in a way that the data is no way recoverable. That is usually accomplished by deliberately overwriting over the data. That is generally the only way to permanently erase data from the hard drive while still retaining its usability.

One can of course break the whole hard drive and its spinning disks with a hammer. That way data is definitely not recoverable, but that also renders that hard disk totally unusable. Another a bit more sophisticated method is to use a degausser. That means using a powerful magnet to completely eliminate drives magnetic field, thus also erasing the data it contains. (Data Security Inc, 2019)

Unfortunately, after this kind of treatment also the low-level firmware required for hard drives use is also lost. As a consequence, it cannot be used again and can only be recycled for its precious metals.

So, as the hard drive itself does not delete data anywhere during, for example, emptying of recycle bin. Instead, the user must overwrite all the data he or she wants gone. This usually does require certain actions with a help of a third-party application.

2.5 Standards

There are some applications to choose from when either erasing individual files or whole hard drives. Later in this paper the focus is on individual file erasure using a certain software. The erasure itself is usually performed by using standards different organizations have established.

One of the simpler standards is HMG Infosec Standard 5. It has two different levels and with base level, erasure is done by only overwriting the data with zeroes once. (Jefcoat, 2017) Depending on how much data one is overwriting, it naturally also dictates the time needed for the erasure. Other factors include how fast the drive is, what kind of interface it uses for connection and the specifications of the actual computer doing the erasure.

The same standard also has an upper level where you do the overwriting three times. First by overwriting the data with ones, then zeroes and lastly ones and zeroes randomly. As the whole data needs to be now overwritten three times, it naturally takes three times as long.

It is debatable if a lot of overwriting rounds are needed for complete data erasure. After even one overwriting round, the data once stored in a magnetic domain gets re-magnetized. Partial restore might be possible in a laboratory environment with the use of equipment like Magnetic Force Microscope, but recovery will be minimal. (Rikhi, 2019)

There are also algorithms which take overwriting to the next level by introducing even more overwriting rounds, like the Gutmann method. It introduces 35 times overwrite, featuring overwriting with zeroes, ones and random data systematically. (Gutmann, 1996) Overwriting a normal hard drive with this method can take up to a week, if it can survive the process without too many errors.

3 Preparing for the assignment

The assignment I got was to test this new filesystem that Microsoft introduced with the company's product to evaluate if the software in question works as intended in this new environment. If it works as intended like in other filesystems, then the result is fine, but if there are complications, my assignment was to find them and potentially find a way to fix them.

The required hardware, software and a place to work in were provided to me. Although, since we were dealing with virtual machines, it was also possible to work on this project remotely. Help was also provided to me by the Blancco Research team.

All the tests were done during summer 2019, and during that time I was also employed by the company. The whole project was also monitored and documented into Jira and Confluence information systems. After getting satisfactory results from the test, I also presented them in a meeting to the stakeholders in Blancco Research and Development teams.

3.1 Introduction to Microsoft Resilient Filesystem (ReFS)

Filesystems

Simply explained, filesystems in computing define how data is stored and retrieved from a storage medium. If there is no filesystem in place, the data in storage just becomes a large data dump where you cannot tell where information is located, where it starts or where it ends.

Filesystem's role is to isolate the data and separate it into pieces where it can easily be identified. There are different kind of filesystems in computing and each has its own structure and logic, but most common nowadays is NTFS (New Technology File System) as it is the default filesystem in Windows NT family. (Custer, 1994) NTFS has many technical improvements compared to its predecessors FAT (File Allocation Table) and HPFS (High Performance File

System) like improved support for metadata, improved performance and reliability and more efficient disk space use.

Microsoft Resilient Filesystem (ReFS)

Microsoft Resilient Filesystem is the new filesystem introduced by Microsoft with Windows Server 2012. The intention was that it would become the "next generation" filesystem after NTFS. It introduces improved reliability and built-in resilience which it accomplishes with checksums created in file's metadata. It automatically does integrity check-ups to prevent files being corrupted and if it finds corruption, it attempts to fix it automatically. Because of this, it is not necessary to manually run Chkdsk like in NTFS filesystem when checking health of a volume.

If you launch Chkdsk on a volume formatted to ReFS, you get the following message:

“The type of the filesystem is ReFS. - The ReFS file system does not need to be checked.”

According to its design, ReFS checks and auto-corrects data on its own. One way to initiate integrity check is to open a file. If there is something wrong with the opened file or a folder, ReFS will automatically try to fix it.

“If the correction fails, the damaged part is isolated, theoretically, without affecting the remaining "healthy" part. If you cannot fix the damage by means of the filesystem driver, you need to recover data using ReFS-capable data recovery software.” – ReclaiMe Datarecovery (2019)

This data correction check can also be run manually with PowerShell command. With Windows's task management, it would also be possible to set it up as a scheduled run. This means that you could have the ReFS run the integrity scan once a week. PowerShell script for that will be included later in the tests performed.

It is also designed to provide new features for performance and virtual workloads that are becoming more and more important today. Good examples of these are real time tier optimization and block cloning. Both are meant to make operation with virtual machines faster.

Support for different ReFS versions, by each Windows version

ReFS	Windows Server 2012	Windows 8.1, Server 2012 R2	Windows 10 v1507 – v1607	Windows Server 2016 TP2, TP3	Windows Server 2016 TP4, TP5	Windows Server 2016 RTM	Windows 10 v1703	Windows 10 v1709, Windows Server 1709	Windows 10 v1803 – v1809, Windows Server 2019, 1803 – 1809
1.1	Default	Yes	Yes	Yes	Yes	Yes	Yes	?	?
1.2	Yes	Default	Default	Yes	Yes	Yes	Yes	Yes	Yes
2.0	No	No	No	No	Default	No	No	No	No
3.0	No	No	No	No	No	Yes	Yes	Yes	Yes
3.1	No	No	No	No	No	Default	Yes	Yes	Yes
3.2	No	No	No	No	No	No	Default	Yes	Yes
3.3	No	No	No	No	No	No	No	Default	Yes
3.4	No	No	No	No	No	No	No	No	Default

Table 1. ReFS support by different Windows versions (Source: Windows ReFS versions.en.md, 2019)

Above you can see the version chart and support for different versions of ReFS. Microsoft has already released many versions of it and even briefly introduced it to Windows 10 operating system. That was when the version number was already at 3 but it was curiously removed in Fall 2017 creators update. ReFS volumes created previously in Windows 10 will still work, but you will no longer be able to create new ones.

So currently you can only use it in server operating systems. It is not certain if Microsoft will ever bring it back to Windows 10.

Challenges

File Encryption – If file level encryption is important to you, ReFS cannot do file encryption. As a side note, you can do BitLocker encryption with ReFS, however, this is whole volume encryption and not at the file level.

ReFS integrity is not on by default (Habets, 2017). Not having integrity stream enabled by default is somewhat dangerous as user may have a false sense of security believing their files are rot-proof.

Some file recover software do not recognize ReFS as a proper filesystem and thus do not work. So, ReFS has been out for a while, but still some of the most known data recovery software do not even support it or mention it in their webpages.

This was not really a challenge for this project, but it is worth mentioning that because ReFS does not support hard links, installing software on ReFS volume is not possible. So, for operating systems and other software you will still need NTFS volume.

3.2 Creation of test data

Test data for ReFS testing

In order to do the tests, data was needed to use for the erasure. This could be any data that is easy to put in a ReFS volume and it should also have files of varying sizes. This was done to mimic a simulated server environment.

That would require obtaining data of specific nature, like big server database files and other files like pictures, text files etc. All of these can be found in the Internet. As database files stack overflow's old database files was used that are distributed for free for testing purposes. For pictures, there are some some royalty free, free for commercial use picture banks. Those are picture sets containing hundreds of pictures.

To combine all of these things, a dataset was made solely for this purpose imaginatively named as "Data".

"Data" was the name of the whole folder that was pasted in ReFS volume and used for testing. There were two major sub folders underneath it:

- Latest-versions
- Backup versions

The whole dataset was roughly 12Gb in size. Next there will be an explanation of how the data sets were comprised

Latest versions

Latest-versions folder had most of the data. Like the name implies, it had the latest versions of the current files needed for this imaginary production server. It had two database files:

1GB database backup file

10GB database backup file

Both of these can be used to restore a full database stack overflow had previously used. They are available in z7 (7zip) format.

Including those, there are two video files, six folders of random stock photos. Most of them even have EXIF data in them. That is data that modern cameras store in each photo. Usually they include information such as camera settings at the time of capture. Each photoset has a certain theme.

There are also some miscellaneous files. Files which include:

- txt files
- Photoshop files
- Excel files

The few text files included had varying data in them. This included the Bible in text file format, list of alphabets and some pseudo randomly generated data.

	Osuus	Koko	Kohte...	Tiedo...	Alikan...	Edellinen muutos	Määr...
Latest-versions	89,4%	10,9 Gt	553	524	29	22.5.2019 12.09.18	
StackOverflow2013	87,8%	9,5 Gt	6	6	0	22.5.2019 7.06.19	
StackOverflow-SQL-Server-20...	10,0%	1,1 Gt	2	2	0	22.5.2019 7.06.19	
<Tiedostot>	0,5%	55,2 Mt	2	2	0	22.5.2019 12.03.43	
VID_20190522_150330.mp4	51,6%	28,5 Mt				22.5.2019 12.03.43	A
VID_20190522_150228.mp4	48,4%	26,7 Mt				22.5.2019 12.02.41	A
Christmas-1	0,4%	49,2 Mt	144	141	3	22.5.2019 7.08.18	
Thanksgiving-2	0,3%	37,3 Mt	100	97	3	22.5.2019 7.09.02	
Hanukkah	0,3%	30,1 Mt	85	82	3	22.5.2019 7.08.37	
Misc-1	0,2%	27,6 Mt	46	43	3	22.5.2019 7.08.45	
Halloween-1	0,2%	24,1 Mt	104	101	3	22.5.2019 7.08.28	
New_Years-1	0,1%	11,0 Mt	35	32	3	22.5.2019 7.08.53	
large	0,1%	10,6 Mt	3	3	0	22.5.2019 7.13.22	
E.coli	41,6%	4,4 Mt				26.2.1997 10.12.00	
bible.txt	36,3%	3,9 Mt				21.3.1997 14.31.58	
world192.txt	22,2%	2,4 Mt				18.3.1997 14.13.34	
cantrby	0,0%	2,7 Mt	11	11	0	22.5.2019 7.15.52	
artificl	0,0%	293,0 Kt	4	4	0	22.5.2019 7.13.16	
aaa.txt	33,3%	97,7 Kt				20.12.2000 9.55.02	
alphabet.txt	33,3%	97,7 Kt				20.12.2000 9.55.04	
random.txt	33,3%	97,7 Kt				20.12.2000 9.55.04	
a.txt	0,0%	1 Tavua				20.12.2000 9.55.02	

Figure 2. Files of the latest-versions

Backup versions

Backup versions is like a stripped version of Latest versions. It included same files, but not all of them. Most notably it only included the 1GB database file.

It also had fewer photosets but the same number of videos. The same random data text files, excel and photoshop files.

	Osuus	Koko	Kohte...	Tiedo...	Alikan...	Edellinen muutos	Määr...
Backup versions	10,6%	1,3 Gt	409	389	20	22.5.2019 12.09.12	
StackOverflow-SQL-Server-20...	84,8%	1,1 Gt	2	2	0	22.5.2019 7.06.42	
<Tiedostot>	4,2%	55,2 Mt	2	2	0	22.5.2019 12.03.43	
Christmas-1	3,7%	49,2 Mt	144	141	3	22.5.2019 7.06.58	
Hanukkah	2,3%	30,1 Mt	85	82	3	22.5.2019 7.07.27	
Misc-1	2,1%	27,6 Mt	46	43	3	22.5.2019 7.07.35	
Halloween-1	1,8%	24,1 Mt	104	101	3	22.5.2019 7.07.13	
large	0,8%	10,6 Mt	3	3	0	22.5.2019 7.12.08	
E.coli	41,6%	4,4 Mt				26.2.1997 10.12.00	
bible.txt	36,3%	3,9 Mt				21.3.1997 14.31.58	
world192.txt	22,2%	2,4 Mt				18.3.1997 14.13.34	
cantrby	0,2%	2,7 Mt	11	11	0	22.5.2019 7.16.28	
artificl	0,0%	293,0 Kt	4	4	0	22.5.2019 7.11.59	
aaa.txt	33,3%	97,7 Kt				20.12.2000 9.55.02	
alphabet.txt	33,3%	97,7 Kt				20.12.2000 9.55.04	
random.txt	33,3%	97,7 Kt				20.12.2000 9.55.04	
a.txt	0,0%	1 Tavua				20.12.2000 9.55.02	

Figure 3. Files of the backup versions

Test data used for this can be found here:

Pictures: (free stock photos)

https://offers.hubspot.com/free-holiday-stock-photos?_ga=2.219242207.711070943.1553693911-1181524008.1553693911

Database files: (stack overflow)

<https://www.brentozar.com/archive/2015/10/how-to-download-the-stack-overflow-database-via-bittorrent/>

3.3 Setting up a test environment for ReFS

In this document, the test setup was made using virtual machines and VMware Workstation. VMware Workstation is a hosted hypervisor which runs on x64 versions of Windows and Linux operating systems. It enables setting up virtual machines on the host machine and using them with the actual machine. Hence, it can be installed on top of an existing operating system like Windows. The other way to do this is with VMware ESXi, which in turn is not a software application, but includes operating system components, such as a kernel. These two are called type 1 hypervisor and type 2 hypervisor, ESXi being a type 1 hypervisor.

The reason for using type 2 hypervisor in this case is that you can also use the actual machine for testing and not just the virtual one. So, first thing was to install Windows Server 2016 on the actual machine, then install VMware Workstation hypervisor on top of that, since that enabled us to also create a virtual Windows Server 2016 as another test environment.

It is easier to do the first tests on a virtual setup as you can take snapshots and revert the machine back to a previous state after a certain test.

3.3.1 Test setup requirements (SW/HW)

The hardware and software requirements for the project are as follows:

HW:

- Workstation PC with decent enough specs (CPU, RAM)
- 2x HDD drives, 1x SSD drives
- Keyboard, mouse and monitor

The two HDDs in this case were Western Digital WD2500AAKX 3,5" SATA HDD and the SSD was Intel SSD 750 series using PCIe NVMe 3.0 x4 interface.

SW:

- Windows server 2016 license
- VMware Workstation 15 license
- File Erasure software
- File recovery software
 - o Disk Drill
 - o EaseUS
 - o RecoveIT

In this setup we did not create any kind of RAID configuration. SSD was used to install an operating system and hypervisor. HDDs were partitioned for ReFS use in both actual machine and virtual.

It is also important to note that if you are using software to recover any lost data, that you install the software before starting the testing. That is because if you start testing, delete a file and then install the software in the same drive, you risk overwriting the data you are recovering.

In this case though, the software will be installed to a different partition than where the tests are done. That is because software cannot be installed to a ReFS filesystem. The file overwriting cannot happen in this case.

For these tests, three different recovery software were chosen. That is because there was no guarantee which would work with ReFS. Some software stated that it does not support ReFS, some said that they do support, and some did not mention it at all.

3.3.2 Creating a bootable USB with Windows Server 2016 image

To create a bootable USB with a specific image, you need to have a software for it. There are multiple solutions for it, but in this case, Rufus was chosen.

Using Rufus is simple and all you have to do is to select the correct USB device up top and then select the .iso image below it. The software itself will recognize the image and select the recommended settings. After creation of the bootable USB, it even reminds you to disable “Secure boot” from the machine you are using the USB stick on. That is because if you are using a machine that has UEFI, it might also be using Microsoft’s security feature “Secure boot”. In order to boot by using secure boot, you would need to have Microsoft sign the boot files.

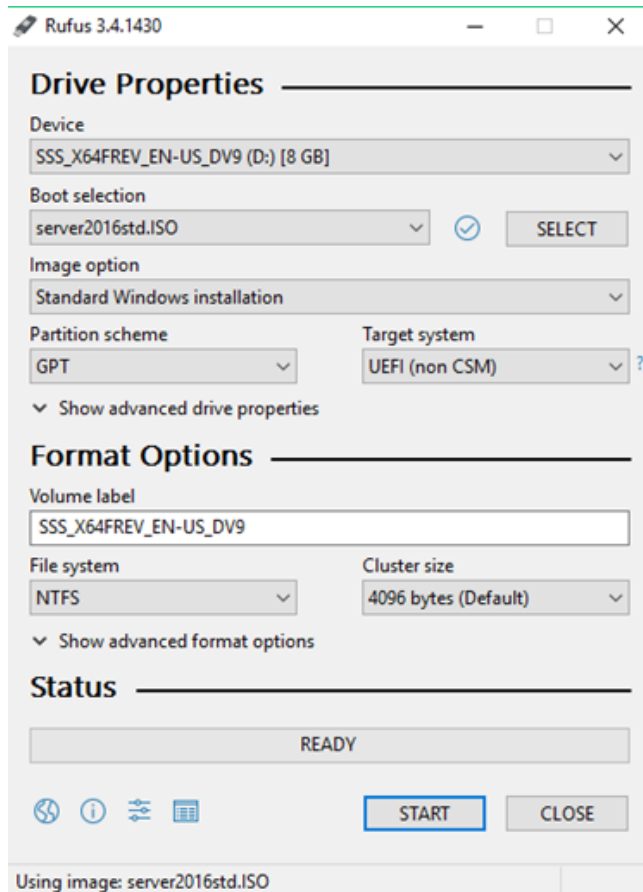


Figure 4. Picture of the USB creation with Rufus

3.3.3 Windows Server Setup

When installing operating system to a new machine, normally a good way to make sure that the installation and all the necessary files go to the correct drive is to disconnect all other drives from the system. That is what is recommended to do as the first step. So, only SSD was left connected and both HDDs were disconnected.

Windows server setup starts like any other operating system setup. Attach the installation media where you have the operating system .iso file to the machine and boot it from that .iso file. In this case, bootable Windows Server 2016 image was created with Rufus. Briefly explained it is a free open source software that can be used to create a bootable USB flash drive. It accomplishes that by

including a bootloader that tells the hardware how to load the installation package in question. This time it is the Windows server .iso.

Operating system installation, which in this case is Windows Server 2016 will start, and you can begin by selecting default keyboard layout, language and time and currency format. After that you can continue by selecting “Install now”.

Next, the version of the operating system to install was chosen. Depending on the installation media used, you might have different operating systems you can install. The main point here is to make sure you install “Desktop experience”. Otherwise you will only end up with command line, without any GUI to help. This is called Server Core and it is a minimal server installation.

In the installation media used, there are two other options – standard and datacenter. For this test environment, standard edition is enough as datacenter edition is more optimized for larger-scale virtualization. With it you can run unlimited number of Windows Server instances which in this case is not needed. (Microsoft Docs, 2019)

Next, you need to accept license terms, then select which type of installation you want. Select custom, and then you need to choose where to install Windows. If you only have SSD connected, like it this case, you should only see the SSD in that list. Select it and then press next.

Windows starts the installation and restarts a few times. Lastly you need to select a password for the administrator user.

Then you should make sure that all the drives you have are seen by the system. If you do not have all drives connected to the machine, shut it down and then connect the drives. Boot it back up and go to disk management. If the disks are not initialized, you might not see them from file explorer normally, but they should be seen in disk management. If the disks are not initialized, you will get prompted to initialize them. Afterwards they will appear on the list and by right-clicking them you can, for example, create volumes.

Next, Windows server operating system and all the applications need to be updated.

3.3.4 Installing VMware Workstation

It is also possible to use VMware Workstation player, but in this example, VMware Workstation Pro was used as it offers more useful features like snapshots. By utilizing snapshots, one can roll back any changes they do. That is useful in case there is a mistake or there is a need to redo a test.

The installation begins by accepting terms and conditions and selecting the installation folder. Installation should be done to the SSD which in this case is the C: drive. When selecting the installation location, you are also asked if you want to install “Enhanced Keyboard Driver”. It will include security improvements and provide better handling of international keyboards, so its recommended to select it. (VMware Documentation, 2019)

Next, you can select user experience settings like product updates on start-up or if you want to join the VMware Customer experience program. You can select them if you want, but the machine that was used did not have internet connection, so these settings were disabled.

Finally, you are asked if you want to create shortcuts. Then the installation begins. If you selected the enhanced keyboard drivers, you will have to finish the installation by restarting your machine.



Figure 5. Starting trial versions for 30 days

The VMware Workstation Pro can be used as a trial version at first, but eventually you will have to provide a license key. On the first start-up you have to either provide a license code or start your 30-day evaluation.

3.3.5 Setting up virtual test environment for ReFS



Figure 6. Creating a new virtual machine

New virtual machine installation can be started from the “Home” screen. Select “New virtual machine” and the installation wizard will open. Select “typical” as a configuration.

In the next step, you need to select the disc image which contains the operating system. In this case, we will select the Server 2016 image. After the image is selected, the installer tells that it will use “Easy install”. This means that the installation process is taken care of by VMware and to initiate it, you will only need to provide info, like windows product key, version of Windows to install and in this case, administrator password.

After that, you need to choose a name for the virtual machine and where it is located. In this the HDD that was attached to the machine was chosen. More

virtual disks were added later to the machine, so it did not matter where the virtual disk containing the operating system was located.

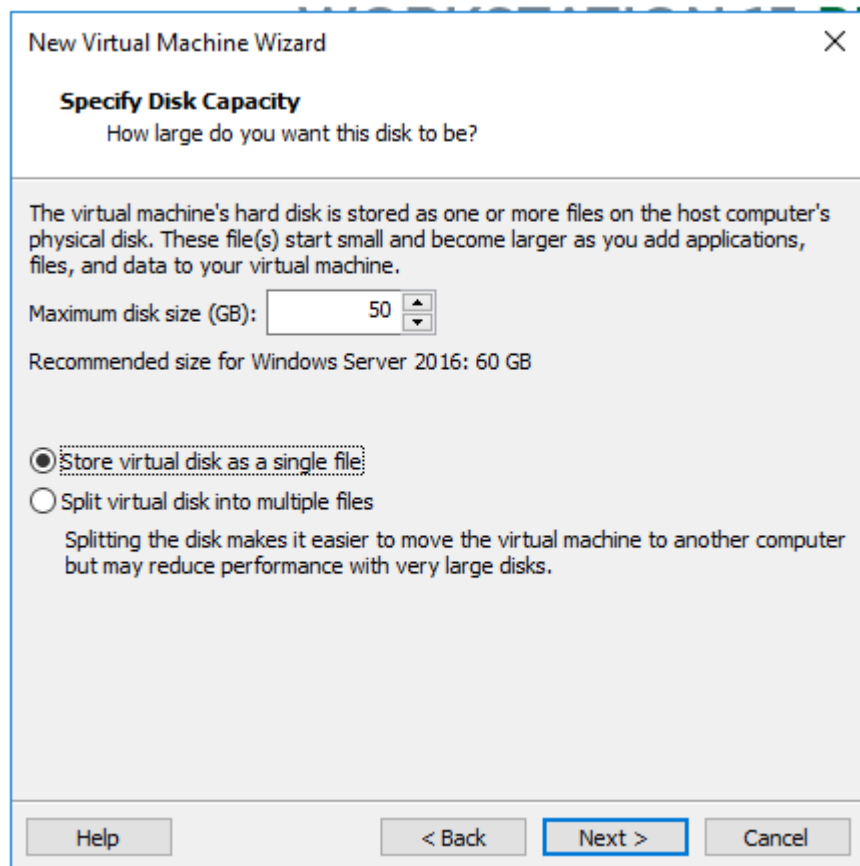


Figure 7. Specifying disk capacity

Then you need to specify the disk used by the virtual machine. The size of the disk created was 50GB. The recommendation can be seen in the picture but making it bigger than 50GB served little purpose for these tests. Also, virtual disk was chosen to be stored in single file

In the last screen VMware is ready to create the virtual machine and you can see the specifications. By clicking "Customize hardware", you can still edit them. The machine was selected to have at least two CPU cores and RAM was increased to 8GB. After that you can click "Finish". If you checked the box to power the virtual machine after creating, it will be powered up and automatically start the installation of the operating system.

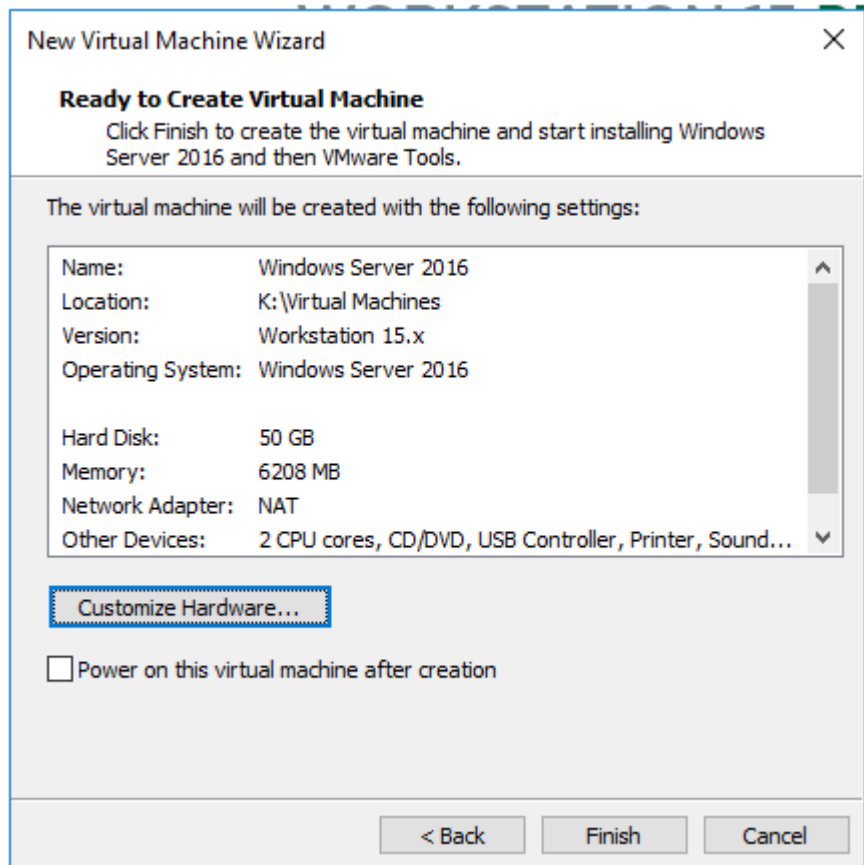


Figure 8. Confirmation page for new virtual machine

More hard disks were needed in this test, so new virtual hard drive was added to the machine next. To do that, you must “Edit the virtual machine settings”. Then on the “Hardware” tab below you can see the button “Add”.

Select hard disk. As the disk type we will need to specify if it is SCSI, SATA etc. As there was a need to create multiple disks with different types, a SCSI disk was added first since SAS (Serial Attached SCSI) disks are one of the most popular in server environments as that is something we wanted to mimic.

Next you need to specify the size. It was made 20GB in size since the tests do not really require too large files. Select the box to allocate disk space now and store the virtual disk as a single file. You will need at least three unused disks to create a virtual disk with parity. Simply explained parity means distributing data among the drives.

The disks you add might be offline at first when you add them to the virtual machine and boot it up. If that is the case, you need to go to disk management and check them from the list below. If it says “Offline”, right-click the disk and select online. Do that to all of the disks there, then close the disk management window and open it again. You should then be greeted with “Initialize disk”-window. There you can select all the disks you want to initialize and choose a partition style. Select GPT as it is the newer one and we do not really need any compatibility with older systems that MBR provides.

Vmware tools

When machine is booted up, VMware tools will be automatically installed to the machine. It provides a lot of utilities which improves the management of virtual machines. It, for example, enables you to be able to drag and drop the File erasure software installation file straight to the virtual machine desktop. So that was done next.

File erasure software

After the file erasure software has been copied to the virtual machine, we can begin the installation. Software installation is straight forward and all you need to do is choose the installation folder. Name of the file erasure software used in these tests is intentionally left out.

Snapshot

After the file eraser is installed, you should create a snapshot of the virtual machine. That will save the current state of the virtual machine and create a point you can return to if needed. It will be useful in case you make a mistake, or just want to test different configurations.

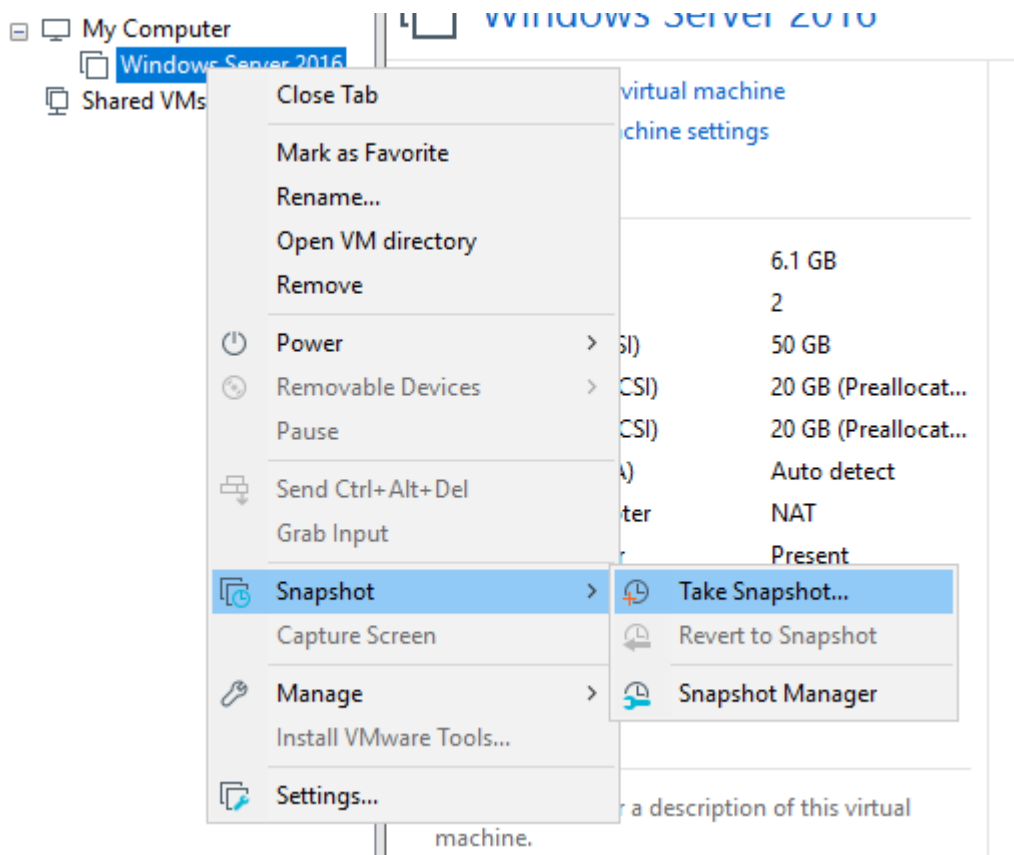


Figure 9. How to take a snapshot

Finally, acquire installation packages for all the recovery software. In this project Disk drill, EaseUS and RecoveIT were used. So, installation of these software should come last. The installation of the three recovery software do not require any specific configuration, and they do not need any licenses since they have free to use versions.

3.3.6 Problems encountered

The first problem that was encountered was that the Serve r2016 could not boot from the machine after it was installed. It was noted that the SSD used during installation was able to create partitions for it and finish the installation, but afterwards the boot was not possible.

The first idea was that the machine was the storage options were incorrect in BIOS settings. SATA emulation was set to IDE, so it was though that it must be AHCI in this case, but it was not possible to choose.

Then it was noticed that the BIOS version was from 2012, which is 7 years old and might not even support PCIe SSDs. So, BIOS was flashed to the latest version, changed the SATA emulation to AHCI since it appeared in BIOS update and then disabled legacy support. After these changes the machine booted to the Windows Server 2016 successfully.

Unspecified disks

If you are having issues, for example, extending disks in an already established storage pool (will be introduced later) it might be because your disks are “Unspecified”- by windows. This can be verified by using this PowerShell command:

```
Get-StoragePool "POOL NAME" | Get-PhysicalDisk | FT  
FriendlyName,Size,MediaType,HealthStatus,OperationalStatus -  
AutoSize
```

With this, you can see a list of all the disks in that named pool. Under “MediaType”, if there are any disks that are “Unspecified” it can cause issues in that storage pool. It was noted that that in all cases the virtual disks in my tests were unspecified.

This can be fixed by manually setting the media type. If all the disks in your pool for example were HDD, you could use this command to set all media types to HDD:

```
Get-StoragePool "POOL NAME" | Get-PhysicalDisk | Where MediaType  
-eq "Unspecified" | Set-PhysicalDisk -MediaType HDD
```

Finally, the plan was to create physical environment too, if virtual tests provided a reason to do more tests. Unfortunately, time allocated for the project ran out.

4 Test scenarios

The main point of these tests is to see how the different features ReFS offers work with secure data erasure. When we are using file erasure software to erase data from ReFS there are two things we want to be sure of.

Any data that is successfully erased from Microsoft Resilient Filesystem with file eraser is NOT recoverable in any way. In this case, we are speaking of “clear” level erasure defined by NIST SP 800-88 (Guidelines for Media Sanitization). Briefly, this means level of erasure that is not recoverable on “keyboard level”. You will not be able to recover any data on a clear level erased storage media without bringing it to a laboratory environment and taking the whole drive apart. Even then it might not be possible. (NIST publication, 2015)

Secondly, file eraser does not interfere with ReFS in any way. This means that it does not, for example, corrupt other files the filesystem has or break the checksums it uses to check the file integrity.

To mimic different setups and ways to implement Microsoft Resilient Filesystem there needs to be different test scenarios to simulate server environments. Tests will be done using both the actual machine and the virtual environment to set them up. Virtual environment will be the first one we test since it is easier to manage and roll back in case of mistakes. More tests will be done on virtual side and only a few of them later in the actual machine.

To do the tests, test data will also need to be created that we are erasing and use data recovery software to try and recover data afterwards. There are many different data recovery software that can be used, so few of the most reliable and reputable were chosen. The test data we are erasing was introduced in section 3.2.

File eraser also has many algorithms to choose from. This means that the erasure can also be done in varying ways. Most notable difference between the different erasure algorithms is the number of overwriting rounds. Other factors include

more complete verification step and the way overwriting is done. Overwriting can be done with either pre-determined data (like write 0's) or aperiodic random data. Multiple overwriting rounds and longer verification increases erasure time.

The chosen erasure standard and verification for this is HMG Lower standard. This is the simplest and probably fastest of the bunch. It will do only one overwriting round and simply do the overwriting with 00's.

Another thing to note about file eraser is that it offers the possibility to erase previous versions. When doing the erasure, it will ask before you start if you want to also erase previous versions of the files. "Previous versions" is a windows feature that that allows you to restore individual files from predetermined restore points. Of course, this can be a security risk, so I will eraser previous versions in all of the test scenarios.

The different test scenarios will be tested as follows:

1. Deploy resilient filesystem on the machine (virtual or actual).
2. Add the test data to the volume that is using ReFS (how the files are placed? Duplicates? Data we want to preserve/data we want to erase)
3. Run file erasure software on the data we want to securely erase
4. After the data erasure is successful, use data recovery software to try and recover any data that was supposedly removed.
5. If data was found, is it usable? This means if we, for example, recover a picture, can you actually open it?

4.1.1 Scenario 1: Testing file integrity and erasure

In the first scenario, file integrity check features that the Resilient Filesystem offers were tested and the erasure in general. Since the ReFS erasure was not tested at all so far, this test also served as a test run for the future cases – is the erasure even possible? This was a legitimate question because some recovery software researched beforehand did not even recognize ReFS as a valid filesystem.

So, the first scenario was done as follows:

- Create Disk pool (from the three disks)
- Create virtual disk with parity
- Create volume > ReFS
- Copy all data to ReFS volume
- Integrity enable > manual integrity check
- Data erasure with file eraser > successful report
- Do manual integrity check
- Data recovery

So, first in order to do virtual disks in Windows server, storage pools needed to be created. This can be done by going to File and Storage and for example trying to create a virtual disk. It will ask to choose a storage pool where we can create one.

For this storage pool, all available disks were selected and the pool was named “Test-pool-ReFS”. This pool was roughly 60GB in size.

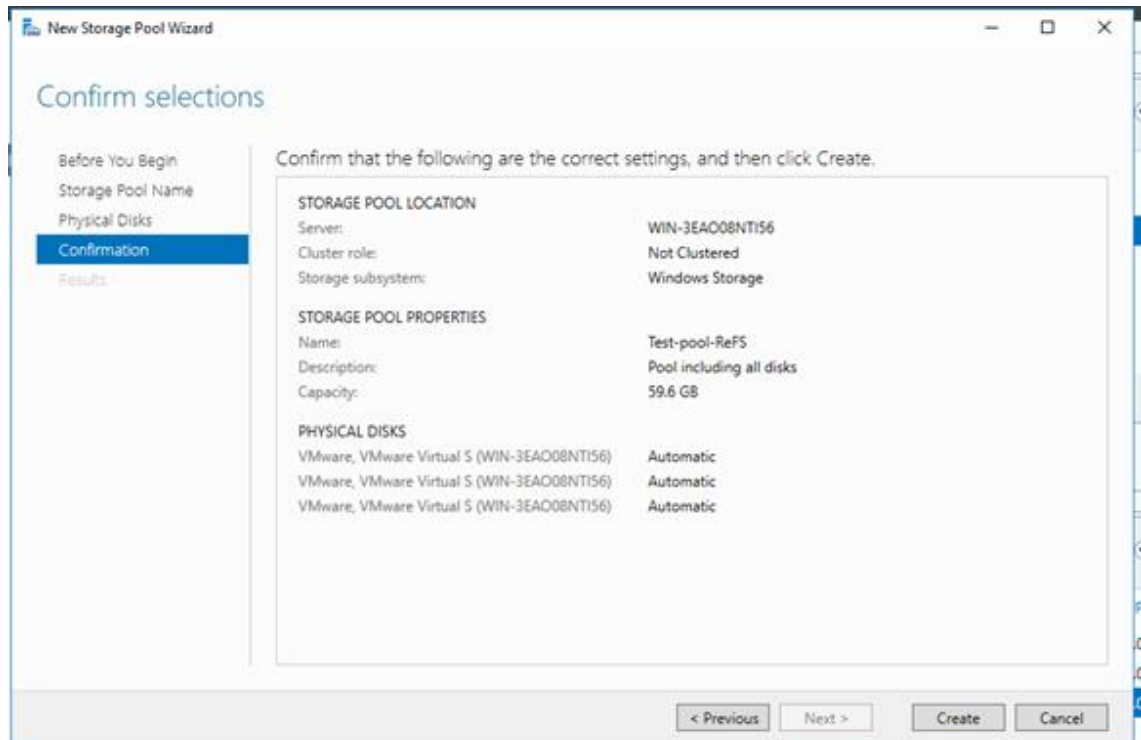


Figure 10. Creation of storage pool

Then, virtual disk needed to be created. This can be done in storage pool tab. There will be a text “To create a virtual disk, start the New Virtual Disk Wizard”.

Start the wizard by clicking it and choose the name and properties for it. For this test, name was chosen “ReFS-disk”, with parity and total size of 36 GB. Test required the maximum size possible with this configuration and in this case, it was 36GB.

Storage tiers were not supported with ReFS, at least during these tests, so selecting that was not recommended.

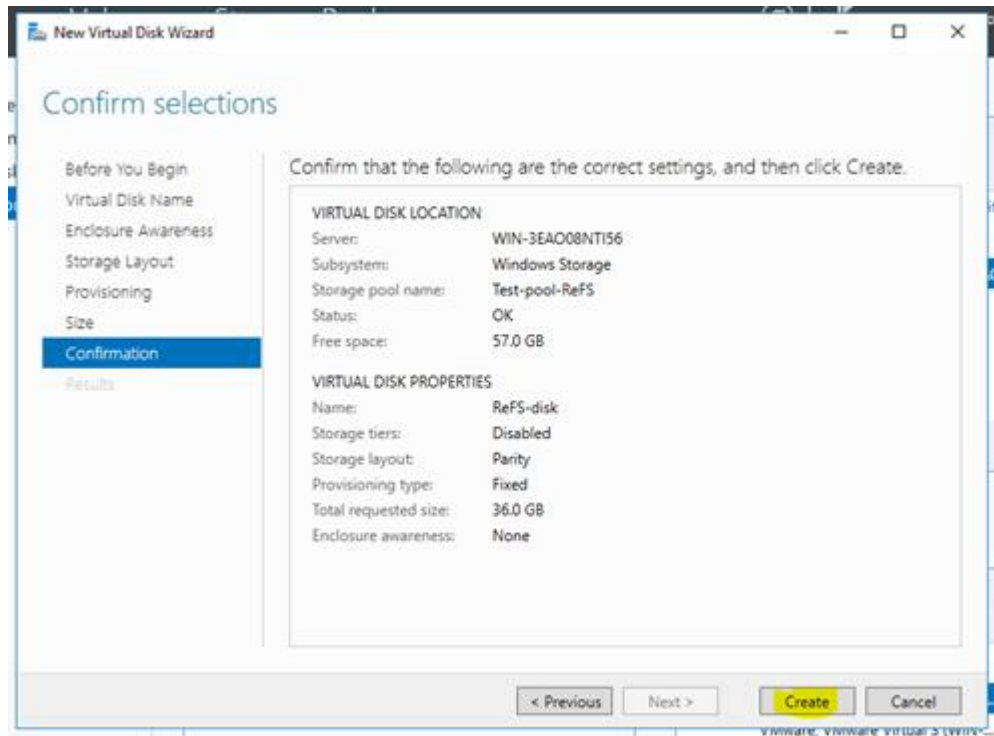


Figure 11. Creation of virtual disk

Lastly, volume for the virtual disk needed to be created and formatted as ReFS. This was done in the same place, but selecting the virtual disk named “ReFS-disk” that was created and clicked “To create a volume, start the New Volume Wizard”.

In the wizard, make sure the correct disk is selected, select maximum size, assign a letter to it and in the File system settings – select ReFS. After confirmation you should have a ReFS drive.




Figure 12. New ReFS volume

Next, all the data used for testing was pasted to the volume. For more information about the data used, go to “Creation of test data”.

File integrity

After the data has been pasted to the drive, integrity checks need to be enabled. This feature is one of the key features of ReFS. It stores a checksum of all the files in their metadata. That means it can run checks to ensure their integrity and possibly fix or remove corrupted files. But unfortunately, this feature is not enabled by default for some reason. In order to enable it, PowerShell needed to be used.

The commands to enable it are fairly simple. To get file integrity, go to the correct directory, and user:

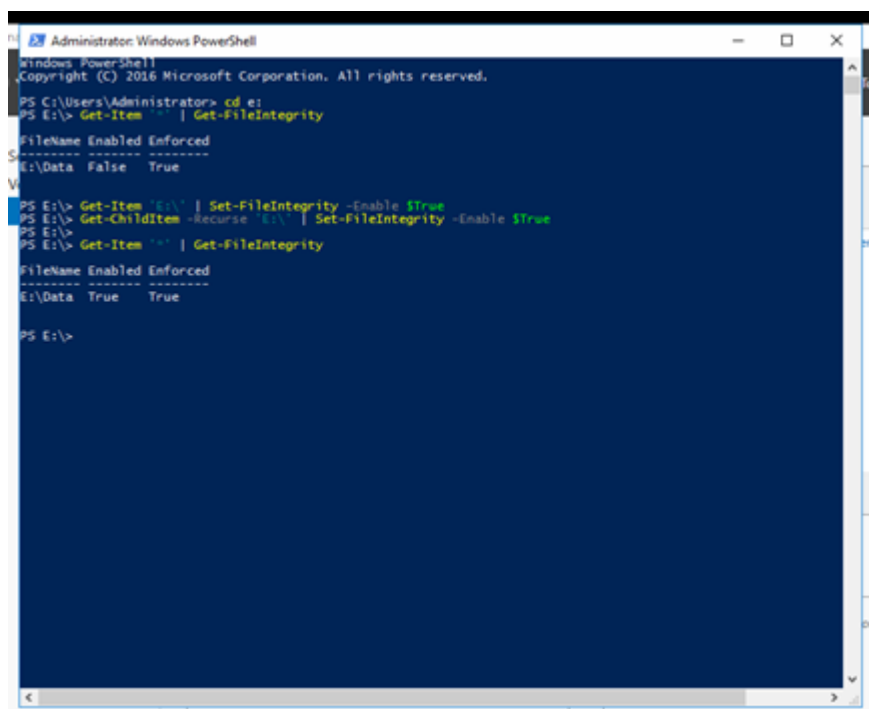
“Get-Item ‘*’ | Get FileIntegrity”

This will list out all directories and tell you if the file integrity is on or off. And to enable file integrity, do:

```
Get-Item 'E:\' | Set-FileIntegrity -Enable $True
```

```
Get-ChildItem -Recurse 'E:\' | SetFileIntegrity -Enable $True
```

In this case the ReFS volume was the letter “E”, so that is where file integrity was enabled. This will enable it to every file currently in E:\ - directory, but also to all the files that are added to it in the future. This means that you do not have to enable it to every individual file afterwards if you add more.



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2016 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> cd e:
PS E:\> Get-Item . | Get-FileIntegrity
-----
FileName Enabled Enforced
-----
E:\Data False True
V

PS E:\> Get-Item 'E:\' | Set-FileIntegrity -Enable $True
PS E:\> Get-ChildItem -Recurse 'E:\' | Set-FileIntegrity -Enable $True
PS E:\>
PS E:\> Get-Item . | Get-FileIntegrity
-----
FileName Enabled Enforced
-----
E:\Data True True

PS E:\>
```

Figure 13. File Integrity enabled

After the commands, you can see that file integrity is enabled.

To run the file integrity check, you can also use PowerShell. There is a task that performs the file integrity check, and we can run it automatically with PowerShell.

The command is:

Start-ScheduledTask -TaskPath "\Microsoft\Windows\Data Integrity Scan\" - TaskName "Data Integrity Scan"

After this, you can check the results in Event Viewer. In this case when scan was executed, it took a few minutes. Results appeared in DataIntegrityScan- folder.

In event viewer when checking the results, you can see multiple "Information"- messages. In this case there was something the check found and successfully fixed. You can see the result here:

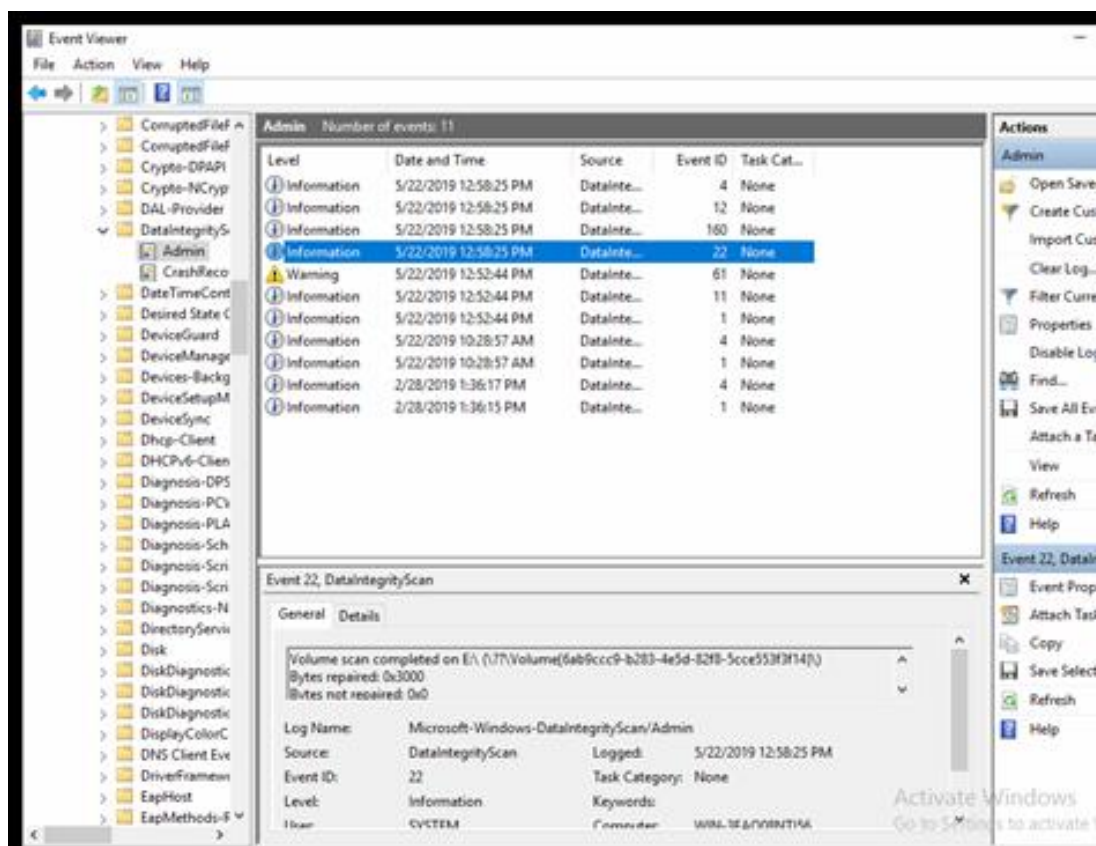


Figure 14. Integrity scan completed

Erasure

The next step is the erasure. File eraser that was already installed on test PC was started and simply "Drag and drop" the "backup-data" folder in the file eraser window started the erasure. This is test was to only erase backup data folder and leave "latest versions" folder intact.

HMG Infosec, Lower Standard was chosen and erasure was started.

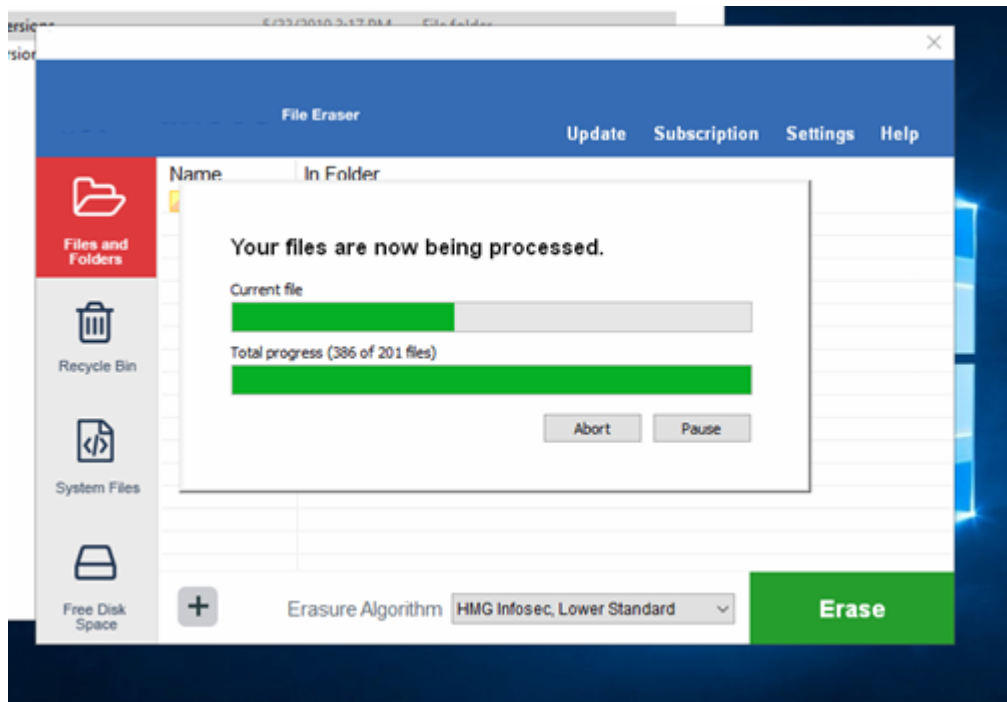


Figure 15. File erasure ongoing

For some reason the files in the “Total progress” seemed to be flipped. As can be seen in figure 15, erasure process was currently already in the file 386 of 201. Anyway, the erasure was successfully completed. Example report of a successful erasure is added as an appendix page.



Figure 16. Erasure successful

After the erasure, to check that file eraser did not mess up the file integrity that was in place, integrity check was run manually. That was done with same command mentioned previously in PowerShell:

Start-ScheduledTask -TaskPath "\\Microsoft\Windows\Data Integrity Scan\" - TaskName "Data Integrity Scan"

Once again it found something to repair, but scan completed successfully after doing the required fixes. Notifications in event viewer don't give out much information beyond this.

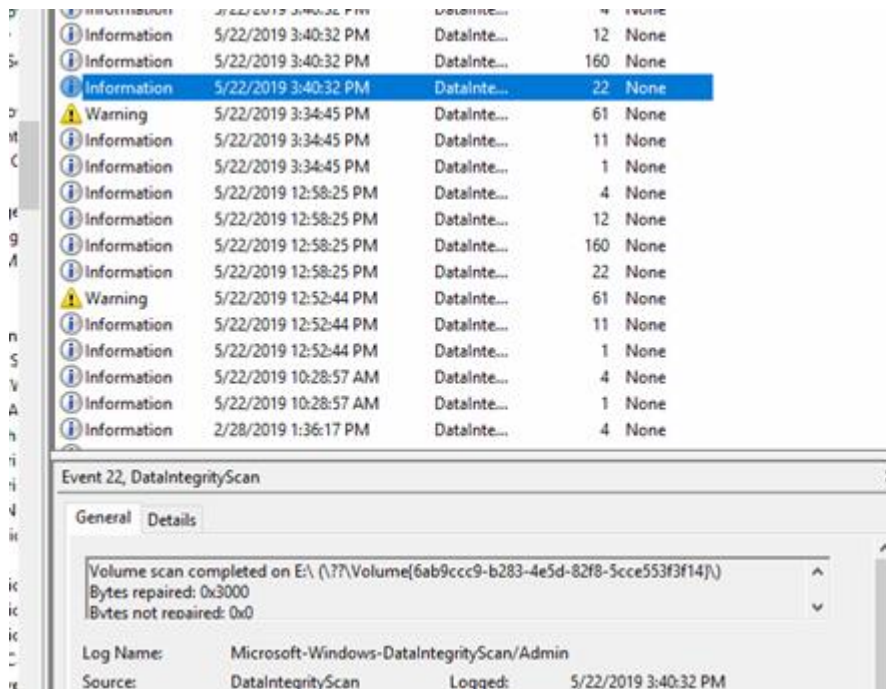


Figure 17. Integrity scan completed

Data recovery

All three data recovery software found recoverable data in this scenario. There were a lot of pictures, same amount of video files than before the erasure and an unusual amount of 7z- files. 7z files were the ones containing the databases used for this. There should only be 10.9 GB of z7 files on the volume, but Disk Drill for example is able to recover 11.71GB worth of z7 files.

Recovery of some pictures and video files was attempted. They were fully watchable without issues. Only their metadata was reset to 1.1.1970. Beginning of Unix epoch time.

Here are all the pictures from the data recovery:

Disk Drill:

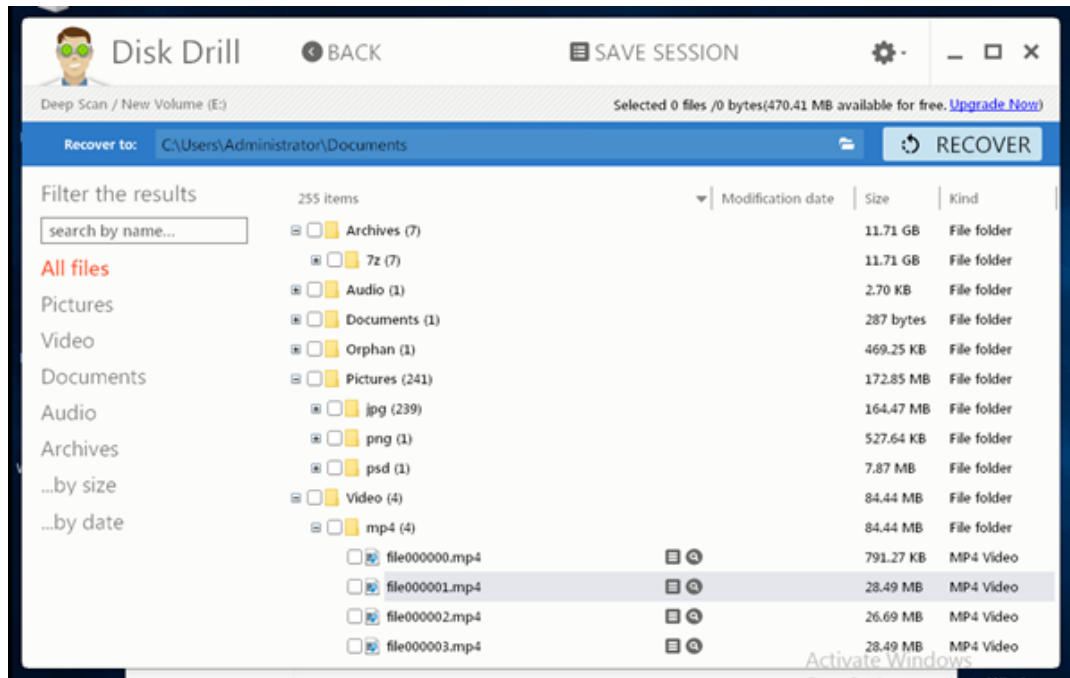


Figure 18. Disk Drill result

EaseUS:

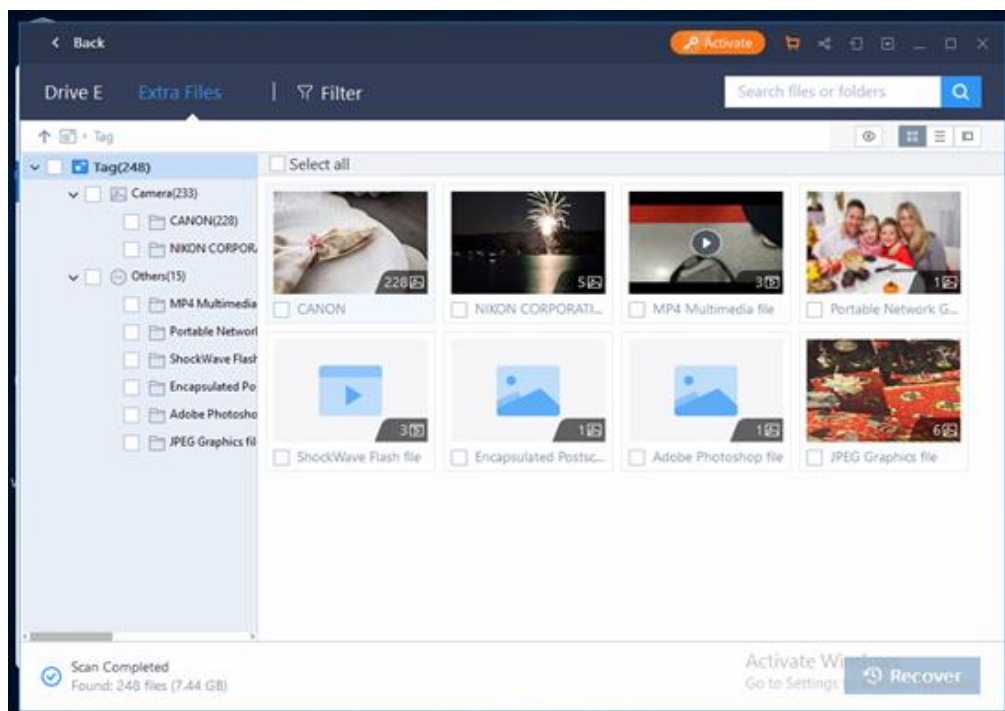


Figure 19. EaseUS result

You can also see that some EXIF data still remain in the pictures. The pictures that have those are grouped in "Canon" folder when using EaseUS

RecoverIT

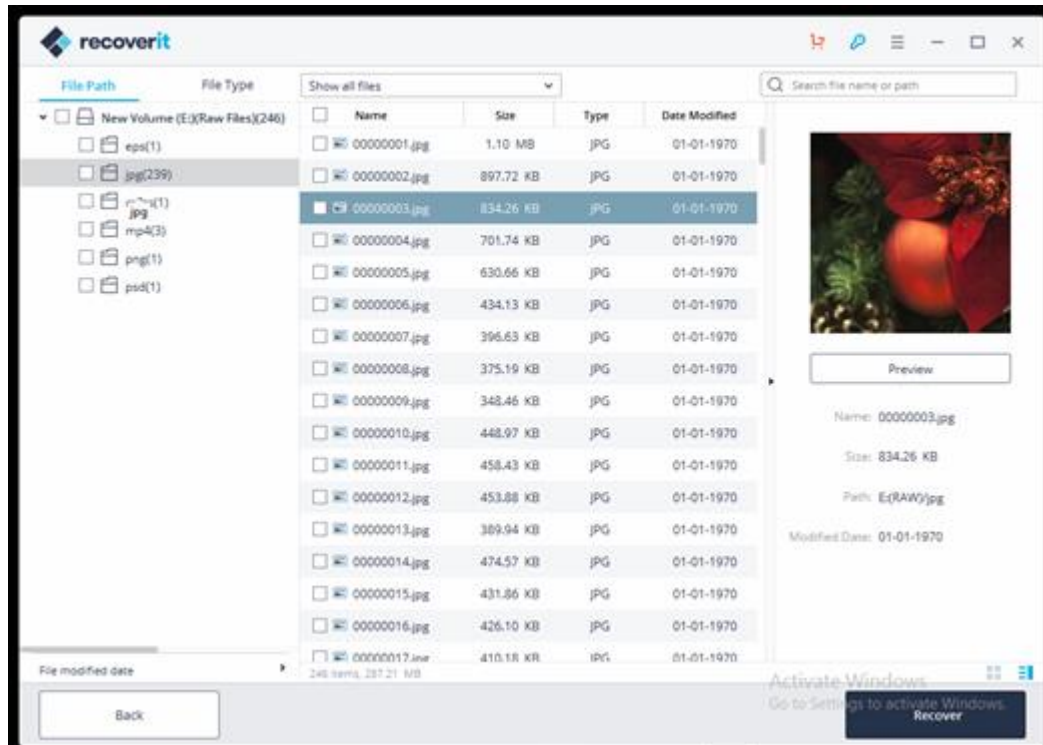


Figure 20 RecoverIT result

4.1.2 Scenario 2: Erasing all the data from ReFS volume

In the previous scenario, all recovery software found recoverable data on ReFS volume even after successful erasure. The next question then was if it was possible that it simply found data in the volume because there was indeed some remaining. Did it recover data that was not even erased or was it just the data that I intentionally left on the volume in other folders?

So next test was to see if the software could recover anything if whole volume was erased. This test was done with two different virtual disks, parity and mirror.

So first it was necessary to delete the disks used in previous test. Virtual machine was reverted to the original snapshot before scenario 1. Then disks were

removed from the virtual machine and the disk files were deleted entirely. There was an issue when too many disks were deleted, which in turn resulted in issues getting the environment running again. So, when deleting virtual disks, remember to check names for the virtual disk files and only delete the files with those names. Also, if you have many snapshots, then names can be very confusing.

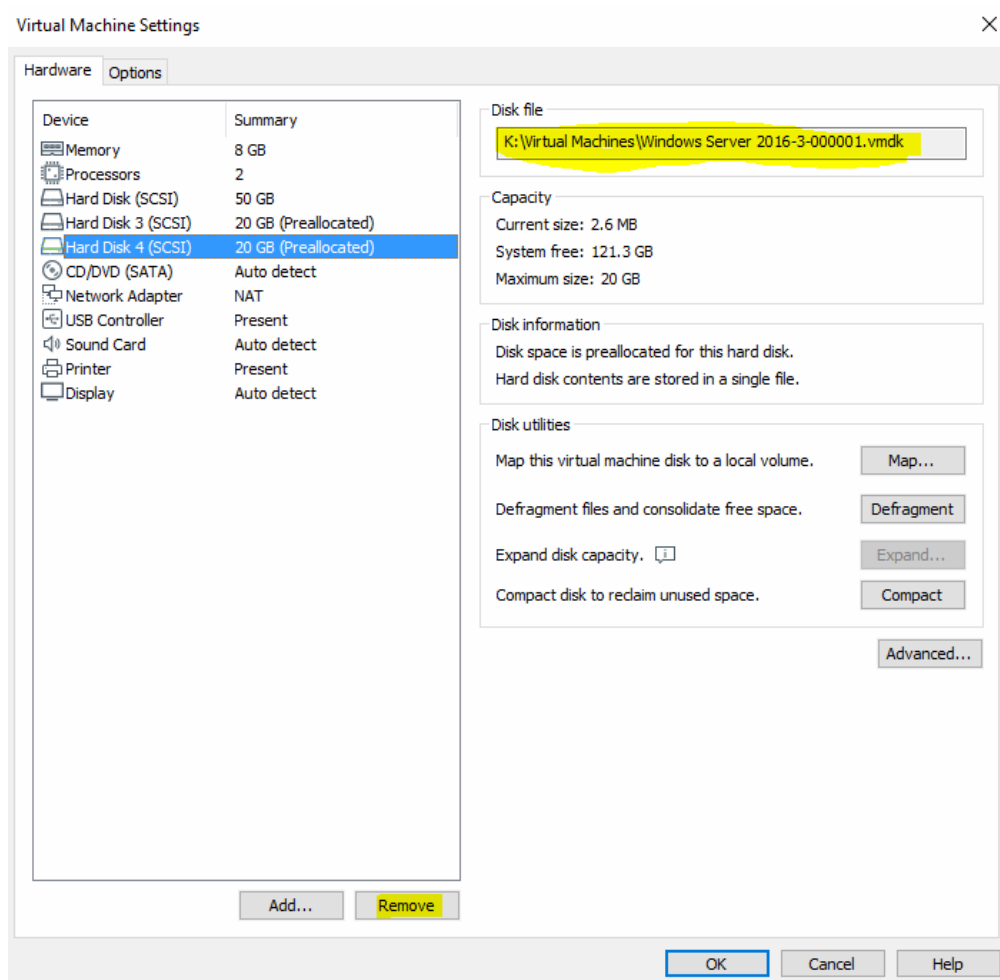


Figure 21. Name of the virtual hard drive

After deleting the files, new virtual disks were created with same sizes and virtual machine was started again.




Figure 22. New virtual hard drive creation

The same steps were performed as in scenario 1, but without data integrity.

- Create Disk pool (from the three disks)
- Create virtual disk with mirror/parity
- Create volume > ReFS
- Copy all data to ReFS volume
- Data erasure with file erasure software > successful report
- Data recovery

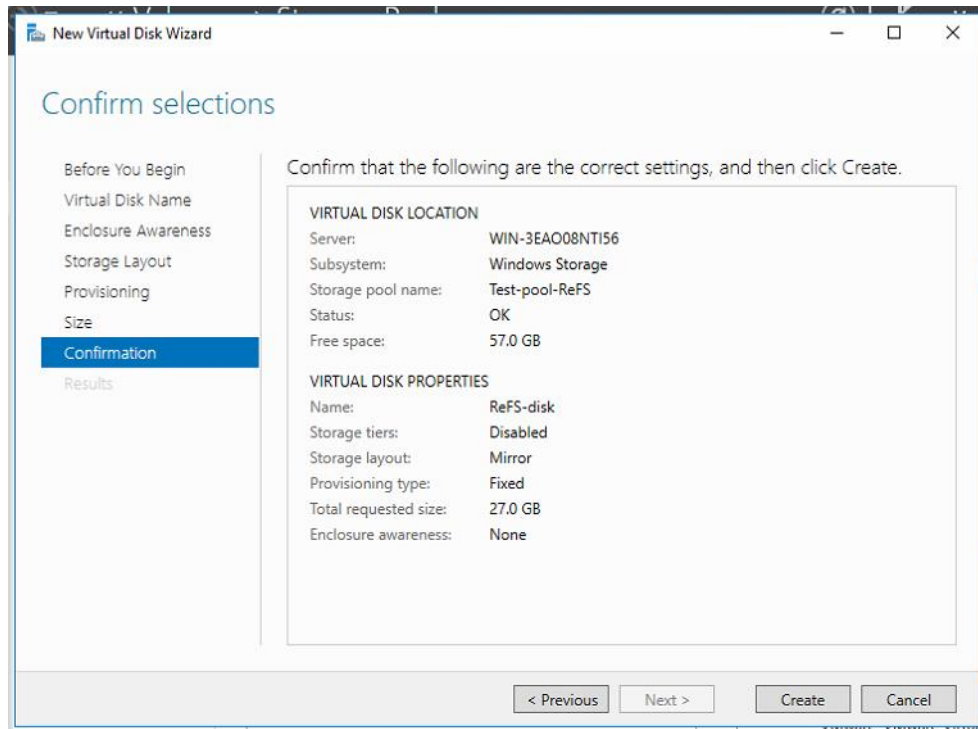


Figure 23. Creation of new virtual hard drive

So, data was pasted again to the ReFS volume, but this time all of it was erased. Nothing was left in the volume by erasing the whole “Data” folder. By doing this it was effectively ensured that the recovery software used cannot recover files that are still remaining in the volume.

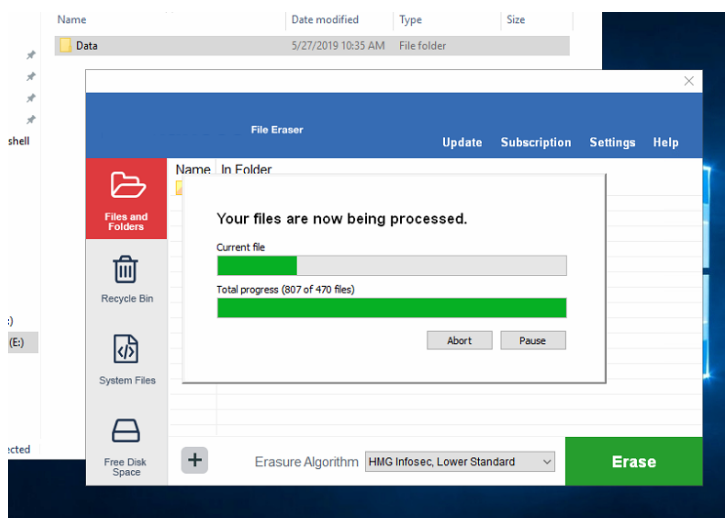


Figure 24. File erasure

After the erasure, none of the recovery software found anything on the volume. All of the chosen recovery software was run with their equivalent of “Deep scan” but they still found nothing. This means file eraser seems to work in this scenario.

Next question then is: What was the data recovered in scenario 1?

4.1.3 Scenario 3: What is the recovered data?

In scenario 3, same setup was used as in scenario 2, since the volume was successfully wiped clean. During this scenario, virtual disk was in mirror. Scenario number three is quite simple, but more difficult to explain.

Since it was noticed that some files were recoverable in scenario 1 and none were recoverable in scenario 2, next answer needed to be found about the data found in scenario 1.

In order to do that, a specific test was devised. The test was done as follows:

- Three random different photos were taken from the data files
- They were put in a folder called folder1
- Copy of that folder was made and renamed folder2
- Now three pictures were in both folders and they were transferred to ReFS volume
- One image in folder1 was opened in picture editing software, and painted a bit over it (picture)
- Then folder1 was erased with file eraser
- Finally, run recovery software to see what version of the picture is the one we can recover?

The edit I did to the picture does not need to be anything specific, as long as it is recognizable from the original.

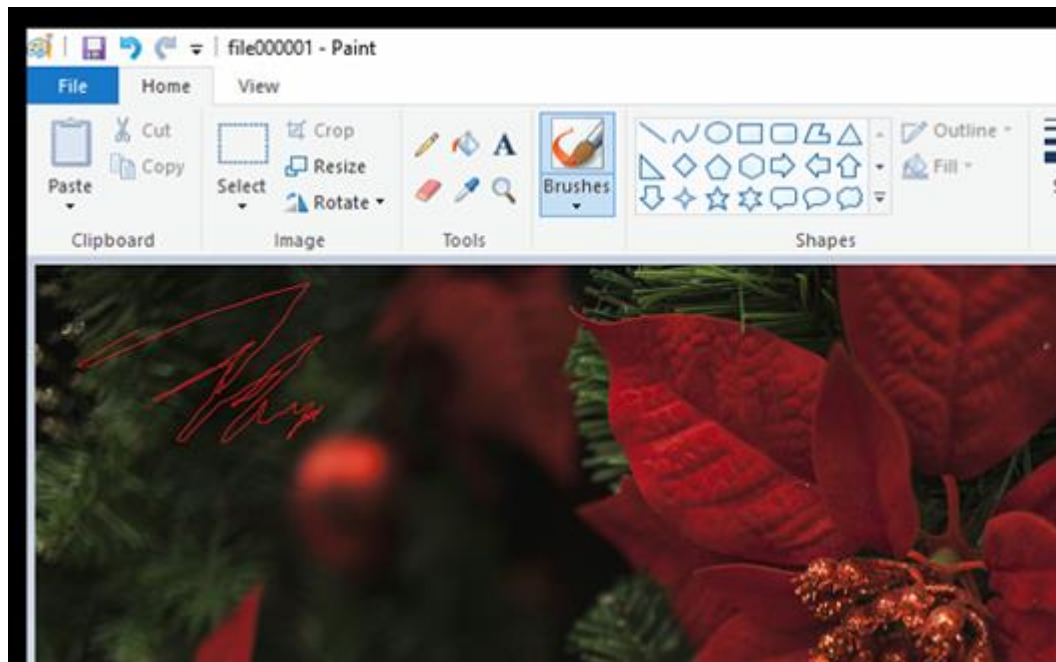


Figure 25. Editing one of the pictures

Then the other folder where the edited picture was located (folder1) was erased with file erasure software and then all three recovery software were run. As a side note, in any of the mentioned scenarios, recovery software were not run at the same time. This had a risk of them interfering each other.

The result was that all three recovery software found four images in the volume. All three pictures that were in the erased folder and a copy of the edited one. So, the edited picture was there twice. But neither of the recovered picture had the edit that was done. This means that file eraser did successfully erase the edited picture, but somehow copy of the picture that was not edited was still recoverable.

4.1.4 Results

	Data found	Data recoverable	Data not recoverable
Scenario 1	X	X	
Scenario 2			X
Scenario 3	X	X	

Table 2. Test results

Above table describes the results from all three different scenarios.

Data found means that recovery software found some data it should not have found after the erasure.

Data recoverable means that the data found, is also recoverable and possible to use or preview.

Data not recoverable means that there were no data during the test that was possible to recover in any way.

4.2 Test conclusion - current hypothesis

There is an issue where some data, or copies of data sill remain on the disk after successful erasure with a file eraser. One explanation for this could be that copy of the erased file was still somewhere located in the ReFS volume.

In scenario 1, there were copies of the same files located in the “Latest versions” folder that was not erased. It could be that because there was a copy of the same file located somewhere else, ReFS was able to recover the deleted one using the recovery software.

In scenario 2, the whole volume was erased leaving nothing behind. So, in the data recovery, we also found nothing to recover. As the volume was completely erased in this scenario, with no copies of the files remaining, not even the recovery software was able to recover anything meaningful.

Scenario 3 is the most interesting one. With a smaller sample size, it was much easier to see what was actually happening during erasure/data recovery. With this test it was clear that data recovery does find something in ReFS volume that it should not have. When erasing three pictures after editing one, if the same picture or maybe even different version of that picture exists somewhere else in that volume, recovery software might be able to recover them. In scenario 3's test, it did not manage to recover all three deleted pictures, but a copy of the edited one. Although it was not the edited one with red paint, it was still the same picture. It could be that ReFS just stored the previous version somewhere during the editing and recovery software was able to find that. But like mentioned before, "Erase previous versions" was selected during all of the scenarios when file eraser asked for it.

This of course proves to be a problem because after the erasure, we get a "Successful" erasure notification and a report even though some data, in some form, seem to be recoverable. This is naturally not desirable.

As the file erasure software in question does not currently promise any support for ReFS, so it is somewhat acceptable. One way to do it now would be to possibly add a note in manual that ReFS is not officially supported. Or to use it at your own risk as the erasure will succeed, but some data may be recoverable.

This should be also properly tested with a physical environment and not just the virtual I had. The plan was to only test physical environment if virtual one provides reason to do so. But unfortunately, time allocated for this project ran out.

The physical test should be conducted with proper server hardware, disk controller and possibly a RAID setup. Maybe RAID 5 and RAID 1+0 with 4 or more disks.

5 Conclusion

Data erasure in this thesis work was done with application running on top of the installed operating system. This of course makes the data erasure a bit trickier as the system's other operations need to be unaffected by the ongoing erasure. These kinds of data erasure software can be used in active datacentres where server cannot be stopped just to erase one drive or volume from a server. But of course, this requires the support of the operating system and filesystem that are currently in use.

Erasure of a single disk can be more straightforward as we can disregard the operating system, filesystem and not worry about the data disk has. It can just be erased in a separate system.

As I worked with a new and not widely in use filesystem in this thesis, it made it much more interesting. With the limited info and other user experiences, it became very hard to predict the results. A lot of the people I spoke with were not familiar with the filesystem nor how it functions. So, it became exiting to see how the tests would succeed.

Personally, I was pleased with the results. If it simply worked without any issues, it would have been a bit more boring to write and come up with different scenarios. I was also glad to be able to provide Blancco with useful data.

I learned a lot about server environments, filesystems, data recovery and erasure during this project. Like mentioned in previous chapter, proper tests with hardware would have been the next step, but unfortunately, I did not have time to properly plan, setup and conduct those tests.

References

Matt Burgess, 2019. What is GDPR? The summary guide to GDPR compliance in the UK.

<https://www.wired.co.uk/article/what-is-gdpr-uk-eu-legislation-compliance-summary-fines-2018> 27.10.2019

Chris Woodford, 2019. Hard drives.

<https://www.explainthatstuff.com/harddrive.html> 27.10.2019

Data Security Inc, 2019. Why use a degausser/hard drive eraser and other FAQs.

<http://datasecurityinc.com/security/degausser.html> 27.10.2019

Katie Jefcoat, 2017. A Comprehensive List of Data Wiping and Erasure Standards.

<https://www.blancco.com/blog-comprehensive-list-data-wiping-erasure-standards/> 27.10.2019

Isha Rikhi, 2019. Why it is Impossible to Recover Data From an Overwritten Hard Drive?

<https://www.stellarinfo.com/blog/why-it-is-impossible-to-recover-data-from-an-overwritten-hard-drive/> 27.10.2019

Peter Gutmann, 1996. Secure Deletion of Data from Magnetic and Solid-State Memory.

https://www.cs.auckland.ac.nz/~pgut001/pubs/secure_del.html
27.10.2019

Helen Custer, 1994. Inside the Windows NT file system.

<https://archive.org/details/insidewindowsntf00cust> 27.10.2019

ReclaiMe Data recovery, 2019. ReFS Recovery.

<https://www.reclaime.com/library/refs-recovery.aspx> 27.10.2019

Oxbadfca11, Windows ReFS versions.en.md.

<https://gist.github.com/Oxbadfca11/da0598e47dd643d933dc>
27.10.2019

Thomas Habets, 2017. ReFS integrity is not on by default.

<https://blog.habets.se/2017/08/ReFS-integrity-is-not-on-by-default.html> 27.10.2019

Microsoft Docs, 2019. Comparison of Standard and Datacenter editions of Windows Server 2019.

<https://docs.microsoft.com/en-us/windows-server/get-started-19/editions-comparison-19> 27.10.2019

VMware Workstation Documentation, 2019. Use the Enhanced Virtual Keyboard Feature in a Virtual Machine.

<https://pubs.vmware.com/workstation-11/index.jsp?topic=%2Fcom.vmware.ws.using.doc%2FGUID-D7E859A1-AD77-41A0-9B20-8B15744056E1.html> 27.10.2019

NIST Special Publication, 2015. Guidelines for Media Sanitization.

<https://www.nist.gov/publications/nist-special-publication-800-88-revision-1-guidelines-media-sanitization> 27.10.2019

Example report from file erasure software

2019-05-22 13:16:53 (+0000),

Data Erasure Report

Erasure Results

Operation - File Shredding

Status: **Successful**

Start/End Time: **2019-05-22 15:16:53 (+0200) / 2019-05-22 15:18:23 (+0200)**

Duration: **00:01:29**

Method: **HMG Infosec Standard 5, Lower Standard**

Target(s): **E:\Data\Backup versions** Size:1378MB **Successful (410 files erased)**

Previous Versions Erased: **Yes**

Software Information

OS Name: **Windows**

OS Version: **10.0.14393**

Computer Name: **WIN-3EAO08NT156**

User Name: **WIN-3EAO08NT156\Administrator**

Report Details

Report UUID: **6db418b1-4465-4cec-aafd-9ecfe1908e06**

Report Date: **2019-05-22 13:16:53 (+0000)**

Software Version: **Blancco File Eraser - Data Center Edition 8.2.1**

Digital Signature: **MEQCIIgfejCk9br255yFAKXrSwHrzVJDbISBcHySwuseLzUeRAiAG5WPaSthsjMe+HI4/x6Ooq2/+OdMiNfrjKyEduFT21Q==**

I hereby state that the data erasure process has been carried out in accordance with the given instructions.

DATA ERASURE OPERATOR

SUPERVISOR