

Long Term Storage in a Surveillance Environment

Anton Danewid `dat14ada@student.lu.se`
Petter Andersson `tna12pan@student.lu.se`

Department of Electrical and Information Technology
Lund University

Supervisor: Maria Kihl & Lars Larsson

Examiner: Christian Nyberg

June 6, 2019

© 2019
Printed in Sweden
Tryckeriet i E-huset, Lund

Abstract

New legal requirements force organizations to increase their data retention of video surveillance for up to a year. This amount of video data requires very large storage capacities, which will drive up the costs for the end user by a large amount. Surveillance data that is old is very rarely accessed, thus there is no need for a short retrieval time for older data. This thesis investigated different storage technologies that could be used in conjunction with a VMS for long term storage of video surveillance.

The use case for this thesis was cannabis cultivation facilities in Canada, which are required to store their surveillance video for up to a year. A technical analysis, as well as an economical one, was made for the different systems, and some were tested in practice. Depending on the amount of data and how long it will be stored, different systems can be chosen. It was concluded that the Veracity Coldstore was a good alternative for the use case presented in this thesis.

Keywords: Long Term Storage, Video surveillance, Cloud storage, Magnetic tape, NAS

Acknowledgements

We would like to thank each of our supervisors at LTH, Maria Kihl and Lars Larsson. Special thanks to Lars, who has helped us immensely with all things regarding the writing of the report as well as always providing tips and moral support, and for always helping us in the never-ending battle against L^AT_EX-formatting.

We would also like to especially thank our supervisor at Axis Communications, Rémy Javelle, for always helping us and answering our questions, whether from Tokyo or Poland.

We would like to thank our technical supervisor Moa Leonhardt for helping us understanding the intricacies of complex software, and we would also like to send our thanks to everyone else that we have been in contact with at Axis, for their very friendly reception and their eagerness to help. We've had a great time working with you all, and we hope to see you again.

Thanks again to Axis, for all the free coffee, vanilla-tea and pastry that supplied us with energy during our endeavor. We would also like to thank Edison Restaurant and Bricks Eatery for their delicious lunches.

We would also like to thank our neighbouring thesis students ("The Stalker Boyz"). We've had some great conversations with these guys, and they helped us realize that even in the worst of times, it could be worse still. We hope that they will one day see the sun again.

Finally, we want to thank Javahuset, also known as Cykelhuset, also known as Tidshuset. The greatest group there ever was. You know who you are.

Popular Science Summary

Video surveillance commonly include many cameras recording in HD, and storing the recordings on a digital storage medium, where the total size of all the recordings can become very large. Today, it is common to store video surveillance for 30-90 days. However, recent changes in the data retention laws – meaning how long the data has to be stored – of some countries, like Canada, means that data now has to be retained for up to 1 year. This data requires a massive amount of storage space, and this amount of storage can easily become very expensive. Furthermore, for certain industries, like the cannabis cultivation industry, failure to comply to these new retention laws can mean an end to the business as it risks losing its license to cultivate cannabis. Therefore, new solutions for storing large amounts of data had to be investigated.

Normally when storing large amounts of data, systems are often configured in something called RAID, which stands for Redundant Array of Independent Disks. RAID is a way for a computer to manage a number of hard drives as one, spreading out information across all of the drives. This increases the speed of the overall system. It also provides redundancy, meaning that the data is stored on several different hard drives; if one drive fails, the data is not lost forever. We investigated one such RAID system, along with other alternative techniques, such as magnetic tape and cloud archival storage. When we had decided which systems were interesting to us, we performed a technical analysis, as well as an economical analysis.

With regards to these analyses, the most prominent solution was found to be a device which can spin down disks that are not being used – in effect rendering them inactive, or idle – as it had a relatively low price as well as very good fault tolerance. Furthermore, it also had very low power consumption, which further reduced the cost of the system along with it being a more sustainable solution. However, this spinning up and down of disks takes time – around 20s – and was considered to be a problem. A solution was designed so that new recordings are saved on the local storage; that is, on the computer on which the recordings are created. This is because, according to Axis Communications, newer recordings are much more often retrieved than older ones, and so they need to be immediately available in the majority of cases. Older recordings, then, are moved to the device in question, which acts as archive storage for the recordings.

Table of Contents

1	Background	1
1.1	Introduction	1
1.2	Background	2
1.3	Long term storage legal requirements	3
1.4	Axis' Requirements for long term storage	4
1.5	Thesis goals	6
2	Different storage technologies and network structures	7
2.1	Hard Disk Drive	7
2.2	Solid State Drive	8
2.3	RAID technology	8
2.4	SAN	12
2.5	NAS	12
2.6	DAS	13
2.7	Magnetic tape	13
2.8	MAID	16
2.9	Cloud Storage	17
2.10	Summary	18
3	Analysis of properties related to storage	19
3.1	Cost of storage	19
3.2	Power consumption	20
3.3	Drive failure & lifetime	21
3.4	Scalability	24
3.5	Performance	24
3.6	Accessibility	26
3.7	Security	26
3.8	Summary of evaluations	27
4	Previous and relevant work for data archiving and data centres	29
4.1	Big data	29
4.2	Environmental impact of data centers	29
4.3	The tiered storage model	30
4.4	Pergamum	30

4.5	Microsoft Pelican	31
4.6	Facebook Open Vault	31
4.7	Summary	32
5	Product presentation and technical evaluation _____	33
5.1	The systems	33
5.2	Technical evaluation	35
5.3	Summary	41
6	Economical evaluations _____	43
6.1	Method	43
6.2	Mathematical model	43
6.3	Price per terabyte	45
6.4	Power consumption	47
6.5	Operation and maintenance	50
6.6	Redundancy and failure	51
6.7	Scalability	53
6.8	Price per TB and power consumption revisited	54
6.9	Evaluation	57
6.10	Discussion	64
6.11	Summary	64
7	Conclusion _____	67
	References _____	69

List of Figures

1.1	The current system design for Axis medium sized systems	2
1.2	A chart over the distribution of retention time settings, here shown in number of days.	3
2.1	A graphical representation of a RAID 6 configuration [10]	10
2.2	A graphical representation of a RAID 10 configuration [2]	11
2.3	Storage servers utilizing RAID technology. [15].	11
2.4	Storage capacity of LTO and other tape technologies. Shows LTO up to LTO-7. Note that LTO-8 is now available in market [19].	14
2.5	A StorageTek tape library system [18].	15
6.1	Price per terabyte, based on product specifications	48
6.2	The power consumption for the evaluated systems. Note the lack of cloud solutions, since the power consumption of cloud-based systems are unknown to the user.	49
6.3	The revisited price per terabyte, with redundancy and scalability factored in	55
6.4	The revisited power consumption, with redundancy and scalability factored in. Again the cloud-based systems are not shown, since the power consumption of those solutions are unknown to the users.	56
6.5	Comparison between different Storinator units and configurations	58
6.6	TCO:s for the Storinator units (in USD)	59
6.7	Comparison between different Coldstore units and configurations	60
6.8	TCO:s for the Coldstore units (in USD)	61
6.9	TCO:s for the cloud service systems (in USD)	62
6.10	The TCO for HPE LTO-8, magnetic tape (in USD)	63

List of Tables

5.1	The measured retrieval times for the different systems	41
6.1	The TCO:s for all of the evaluated systems in descending order of cost. The amount of data to be stored in this example is 500TB and the time-period is 5 years.	64

Acronyms

VMS	Video Management System
FPS	Frames per second
HDD	Hard Disk Drive
SSD	Solid State Drive
RAID	Redundant Array of Independent Disks
SAN	Storage Area Network
NAS	Network Attached Storage
LTO	Linear Tape-Open
LTFS	Linear Tape File Format
MAID	Massive Array of Idle Disks
AWS	Amazon Web Services
MSRP	Manufacturer's Suggested Retail Price
TCO	Total Cost of Ownership
W	Watt
kWh	kilo Watt hours
USD	United States Dollar
bps	Bits per second
Bps	Bytes per second
API	Application Programming Interface
SDK	Software Development Kit

1.1 Introduction

Archiving digital media for the long term is a challenge that many actors are faced with, especially organizations that handle large amounts of data such as Google and Amazon. Data is business for these organizations, and so it is therefore vital that data can be preserved for the long term for backup purposes as well as for business purposes. However, much of the data is very rarely, if ever, accessed, but it still has to be retained for either legal reasons, or again, for business reasons. Therefore, there is ongoing research and development into the field of storing large amounts of data effectively and cheap, while still achieving acceptable retrieval times.

However, the organizations mentioned above are not the only organizations that have to store data for longer periods of time. Today, the average retention time for video surveillance is around 30 to 90 days. Depending on the camera type and amount of video, this data takes up a large amount of space on a storage system, which is why the retention time is set to a limit which might, at a first glance, be considered as short. Nonetheless, this trend is changing as there are currently new regulations in place that affect different sectors in various countries, which force organizations active in these sectors to retain video data from surveillance cameras for longer periods of time.

When storing surveillance video data for longer periods of time, the storage space needed reaches astronomical sizes, and of course, the cost follows, especially for the end user who has to pay for the storage space for all the data. Furthermore, the data might not contribute financially to the overall organization at all. There is therefore a need to investigate different solutions to this problem. One topic is: what can be done with regards to choosing a new storage solution that is relatively cheap, but still fulfills the requirements from both the end user and the surveillance camera vendor?

In the case relevant for this thesis, organizations are by law required to store their surveillance videos for at least 1 year, which is an increase by 4 times of the default retention time currently set in medium sized systems managed by Axis Communications, for which this thesis was written. For a medium sized organization that uses a surveillance system, this results in a massive amount of data that is not directly related to the business of the organization. Many other

organizations that have requirements on data retention, but in different areas that are not video surveillance, will not have the same storage requirements and challenges, since video data needs large amount of storage space. After 1 year, when the retention time is over, the video data can safely be deleted without any repercussions from law enforcement agencies.

1.2 Background

To get a better grasp of the problem, it is important to get an understanding of how a typical video surveillance system currently in use, works.

The standard setup for using Axis cameras in a medium sized system is using the provided Axis Video Management System (VMS). The cameras are connected to a server running the VMS via a network switch. All communication with the cameras is done via the server. Both the VMS client and the server operates on a Windows computer which is connected to the same network. It is through the VMS that the user can change settings such as storage locations, and furthermore manipulate the cameras. The video storage can be on the server itself, or on another storage device such as a NAS or the computer that runs the VMS. The user can specify retention times for specific cameras. When the recordings are outside of the retention time window, they are removed. It is also possible to specify an unlimited retention time, which means that the system will remove the oldest recordings to make space for new ones, when the storage is full.

All of the recorded data is visible inside of the used VMS for viewing. The user can go back to the oldest recordings for playback.

The system setup is illustrated in Figure 1.1.




Figure 1.1: The current system design for Axis medium sized systems

1.3 Long term storage legal requirements

As can be seen in Figure 1.2, the most common retention time for Axis medium sized systems is 30 days, and the second most common retention time is 90 days. This is probably due to the fact that the default retention setting in the VMS is set to 90 days, and in earlier versions of the VMS it was 30 days.




Figure 1.2: A chart over the distribution of retention time settings, here shown in number of days.

The long term storage issue arises mainly from legal requirements in certain countries. These requirements are fairly recent, and thus there is no solution to this problem in place at the moment for Axis. In the Middle East, the retention time for video surveillance is increasing to six months. In Canada, where recreational cannabis is now legal, the cannabis cultivation facilities needs to cover the whole perimeter with cameras. The video has to be saved for a year minimum, in 15 FPS with 720p as the lowest resolution [8, 20]. In the states in the US that has legalized recreational use of cannabis, there are also similar laws of video retention, but they are less strict than their Canadian counterparts.

Since the data is high definition video, it will take up a large amount of space. The data will however rarely be accessed, since it is only kept for legal compliance. According to product managers at Axis, older recorded surveillance video is very rarely accessed – or played back – at all. Therefore, surveillance data like this can be referred to as cold data. The concept of hot and cold data will be described in more depth in Section 4.3. This means that the data retrieval speed will not have to be fast; in fact, several hours is considered acceptable since the data will most likely only have to be retrieved when required by law enforcement agencies or alike. It can therefore be stored on a lower grade storage system that is not as expensive, and does not consume as much power.

So far the Canadian retention laws for cannabis cultivation are the strictest and presents the toughest requirements, and thus, they will be used as the basis requirements for this thesis.

1.3.1 The Canadian retention laws and requirements

The regulations presented in this chapter is a combination of the actual regulations, in combination with an interpretation by Axis, and customer needs and requirements. As mentioned, the Canadian laws state that the surveillance system at a cannabis production facility needs to be set to continuous mode and have a frame rate of at least 15 FPS. In contrast to motion-detect mode, continuous mode saves all the recordings to disk, no matter if a motion or change has been detected in the image. The system coverage is required to encompass 100% of all interior building space, with a minimum pixel density of 140ppf (pixels per foot). The exterior of the building is also required to be 100% covered by the surveillance system with the same pixel density (140ppf). The recording has to be retained for at least one year after it has been made.

Regarding the repercussion for failing to comply with the data retention laws, they are indeed quite severe. Government auditors who inspect the cannabis growing facilities conduct a series of checks at the site. If the facility fails to record for a specified time and date, it is considered a fail on that requirement. If the facility fails to satisfy a minimum number of requirements, they risk losing their licence to cultivate cannabis, resulting in financial disaster and possible bankruptcy.

If data loss were to happen, the cannabis cultivator will most likely complain to the system integrator, who is the one buying the equipment from Axis. There is then a risk that the system integrator drops their contract with Axis, which would mean a financial loss for Axis as well. It is therefore of utmost importance that the recordings are always available, for all parties involved.

In short, the following legal requirements apply:

- 720p resolution.
- 15 FPS.
- The surveillance coverage needs to encompass 100% of all interior space with a minimum pixel density of 140ppf.
- The surveillance coverage needs to encompass 100% of all exterior space with a minimum pixel density of 140ppf.
- Continuous recording, no motion detection mode is allowed.
- Retention time minimum of 1 year.
- Data needs to be fairly accessible.

1.4 Axis' Requirements for long term storage

In addition to the data retention laws, Axis has their own set of requirements for a new storage system both in terms of cost and technical requirements.

1.4.1 General requirements

For long term storage, certain aspects of the current system setup will have to be modified. The two most important requirements are lower cost per terabyte and lower power consumption for a new storage system. These in term sum up to a total cost of ownership (TCO), measured in USD.

As previously mentioned, most recorded video will never be played back. This means that the data does not have to be saved on a high performance storage system. These systems are often very expensive as well as having a high energy expenditure. Therefore it is acceptable if older data is stored on a lower tiered storage system that is slower as well as cheaper.

It is important that the archived data in the long term system is still in the correct format, meaning the file format specific for the VMS used. Since the data format is a proprietary format, it contains metadata such as timestamps, so as to be usable inside the VMS. The data can be exported as standard video formats, but this would mean that the video is not playable in the used VMS anymore since it would lose the metadata. Being playable inside of the VMS refers to actions such as playing multiple recordings synchronized and simultaneously. If the recordings are converted to standard video files, the user would lose this functionality.

The solution also needs to be scalable, meaning that if the storage requirement increases, the underlying solution does not have to be changed in a drastic way.

1.4.2 Technical requirements

As mentioned in the legal requirements section, the resolution required is 720p as well as 15FPS. The cameras need to be in continuous recording mode as well, meaning that they are recording continuously to disk, in contrast to motion detection based recording. With Axis calculations, this amounts up to about 2Mbit/s per camera.

The usual amount of cameras in the real world scenario is in between 10 and 60 cameras. With these properties in mind, the total bitrate for all cameras amounts to between 20 or 120Mbit/s. For a days worth of recording, this can amount up to $\frac{120 \cdot 60 \cdot 60 \cdot 24}{8} = 1296000\text{MB}$ per day, or circa 1.3TB. For a year, this amounts to 474TB. Therefore, the new storage device needs to handle sizes of this magnitude. It also has to be able to handle a bitrate of 120Mbit/s to be able to save the recordings to disk.

However, it is important to note that 2Mbit/s per camera is a conservative number, and that Axis compression technology will help produce a lower bit stream. For the purpose of this thesis, the conservative number was used in calculations as a worst case scenario.

Furthermore, it is also preferable if the new storage solution consumes as little power as possible.

Also, the VMS might have to be changed or have added functionality, to accommodate for this new storage type.

1.4.3 Technical and general requirements summary

In short, the storage system needs to be able to handle these requirements.

- Around 470TB or larger, or the equivalent of 1 year of video recordings from 10-60 cameras.
- Handle bitrate larger than 120Mbit/s, or the equivalent of 60 cameras.
- Save the data in the same format as on the regular current storage alternative, so that it is still viewable in the VMS.
- Consume as little power as possible.
- Be cheaper per TB than current storage solutions.
- It is acceptable if the data retrieval times are in the magnitude of hours.
- The legal requirements imply that redundancy is a very desirable trait for the system to possess.

1.5 Thesis goals

The purpose of this thesis was to investigate the topic of long term storage applied to video surveillance. The main goal was to investigate different storage solutions, mainly from two perspectives: economical and technical. The economical perspective tried to predict how expensive the solution would be, taking different parameters into account. The technical perspective refers to how the system will behave together with the video surveillance system. How could such a solution be used together with the video surveillance system, and how would such an integration look like? This section also included testing of some of the solutions to further investigate them, as well as integration and implementation of some of the systems.

Different storage technologies and network structures

In this chapter, different storage types and technologies, and techniques related to storage, will be presented to get the reader acquainted with some of the concepts that are present in the rest of the thesis.

2.1 Hard Disk Drive

The most common form of data storage is the Hard Disk Drive, or HDD for short. The HDD is electromechanical in its nature, being composed of a magnetic storage and rotating disks which read and write data to the magnetic storage. HDDs utilize random-access, meaning that data can be accessed both sequentially and in any order [55]. Finally, HDDs are *non-volatile*, which means that they retain data stored on them when they are powered off. This is what allows them to be used as what is commonly referred to as *secondary storage*, in most computers, where the *primary storage* (sometimes called working memory), refers to the RAM, the processor registers and the cache [51]. For this reason, HDDs are utilized in - among other structures - the RAID structure, which will be discussed below.

The main benefit of using an HDD is that they are very cheap. Since the technology is so old (it was introduced by IBM in the 1950s), a lot of work has been put into making HDDs as efficient as possible with regards to price per gigabyte, and as such, it is possible to store large amounts of data on a single HDD. However, the HDD is susceptible to a lot of errors. If the HDD is not turned on, the standard retention time ranges from everything between five to twenty years, if given the right conditions with regards to temperature, humidity and overall physical activity surrounding the drive [45, 56]. The deterioration of the data can depend on a lot of different factors, ranging from corruption in the magnetic fields of the storage, or the fact that magnets themselves lose their field of strength as the years go by, to the rotating disks being damaged by physical force, which can result from the HDD getting hit, or simply by it losing power at the wrong time, causing a "hick-up" in the rotation [56].

2.2 Solid State Drive

The Solid State Drive (or SSD for short) is a form of non-volatile memory that is, like the HDD, utilized as secondary memory in some computer architectures. However, they differ from HDDs. The main difference between the SSD and the HDD is that the SSD uses integrated circuit assemblies. What this means is that SSDs store their data in electronic cells (i.e. in transistors), and not in a physical disc like the HDD, and in effect, the SSD has no moving parts. This means that SSDs are more resilient to pure physical force. It should be said that SSDs are, of course, more expensive than HDDs.

Something to consider with SSDs is that since they store data in the form of electric charges, that data will slowly leak out and disappear as time goes on if they are turned off. This is affected by outer factors such as temperature, but the general time before data starts to leak is between one and two years, though some sources claim that leakage can begin after as early as three months [41]. With this in mind, it is important to take into consideration for how long data needs to be stored or archived before being accessed at all.

Another thing to remember when working with SSDs in large-scale operations, like in data centers, is that there is a finite number of write operations allowed to be performed. On a personal computer, this is oftentimes not something to worry about, since the number of write operations is quite high, but when working on industrial-level storage scale this fact should not be dismissed. It is common for manufacturers to refer to this finite number as terabyte written, or TBW. TBW for personal use SSDs (with a size of around 126 to 256 GB) are normally in the range of between 60 to 150 TBW, which means that a very large amount of data has to be written for the limit to be reached. It is worth noting, however, that a test performed by one of Germany's most respected magazines for IT and computers, to see how big the TBW was in practice, showed that almost all of the SSDs tested performed better than what was promised by the manufacturers [39].

TBW is not the only measurement used to determine the endurance of an SSD. Another measurement, that is more common in the industry, is Drive Writes Per Day, or DWPD. Calculating the DWPD is done as follows: take the TBW of the SSD divided by the warranty in days times the size of the SSD. An example, having an SSD with a warranty of 2 years, a size of 200 GB and a TBW of 1000: $1000 \text{ TB} / (365 * 2 * 200 \text{ GB}) = 6.85 \text{ DWPD}$, which means that this particular SSD can have its entire storage space refilled almost seven times a day for the two year period. There is currently an ongoing discussion as to which measurement is most accurate, and relevant, but it is often an advantage to take both into consideration [31, 35]. One last thing to consider with SSDs is that it is harder to recover lost data from a failed drive, when compared to the traditional HDD, since it is harder to get access to the storage device due to how the hardware is constructed, than it is to repair the magnetic storage disks of the HDD [46].

2.3 RAID technology

RAID stands for Redundant Array of Independent Disks, and it is a technique utilized to increase performance of the disks and/or increase redundancy within

them, as a way to safeguard against disk failure. The concept was introduced in 1988 by David A. Patterson, Garth A. Gibson and Randy H. Katz at the University of California, Berkeley [50]. With RAID technology, data is distributed among several drives in parallel, in different ways, called RAID levels. Each level has a specific goal, such as availability, capacity, performance and so on.

RAID can be implemented both on a hardware and on a software level. The unit responsible for managing the RAID system is called a RAID controller, and this can be either a piece of hardware or a software program. The controller is used to offer a different layer of abstraction, so as to be able to view the drives in the array as logical units with regards to the operating system, to allow the user to define specific schemes which allow for data protection.

As mentioned, there are several different RAID levels. Here, the most common ones will be briefly described. Specifically, see figures 2.1 and 2.2 for graphical representations of RAID 6 and RAID 10, as these configurations will be important in the upcoming analysis.

2.3.1 RAID 0

RAID level 0 is the most simple level where the main goal is to increase performance. It splits the data across the disks in the array, allowing all disks in the array to function as one single partition. This increases performance by, for example, letting users read data from a file that may be stored partially on several disks, thus allowing the capacity of all the disks to be utilized at once. This level lacks any form of redundancy, though, so a disk failure results in the array getting broken and data to be lost.

2.3.2 RAID 1

RAID 1 has the main function of providing redundancy using a process called "mirroring" - what this means is that data is written to several disks at once, instead of just one. If one disk fails, a "mirror image" of that disk exists on another disk, thus protecting the data. RAID 1 thus requires at least 2 disks.

2.3.3 RAID 5

This RAID level has two functions; it functions both as RAID level 1 - splitting data and thus increasing performance, and RAID level 2, but in a different way. Instead of mirroring the entire disk, it stores what is commonly referred to as parity information. Parity information is a small piece of information that can accurately describe a larger piece of information, and this protects the disks against failure. RAID 5 requires at least 3 disks.

2.3.4 RAID 6

RAID 6 is almost the same as RAID 5, but the difference is that it stores an extra piece of parity information, effectively doubling the protection, since if one parity-containing disk fails, there is another one containing the same information. As with RAID 5, the parity information is simply a smaller piece of information

that describe larger pieces of information, and as stated, all parity information in a RAID 6 array exists at two places at once. What this means is that two disk failures can occur at once, and the data will still be protected. Due to this, RAID 6 requires at least 4 disks.




Figure 2.1: A graphical representation of a RAID 6 configuration [10]

2.3.5 RAID 10

RAID 10 is a combination of levels 1 and 0. It is often called a *hybrid* RAID. What it does is that it creates a mirror (RAID 1) of the data that is split between the disks (RAID 0). It requires a higher number of minimum disks than the individual levels 0 and 1, but is useful since it provides both a higher performance, gained from level 0, and a higher degree of safety, gained from level 1. Like RAID 6, RAID 10 requires at least 4 disks.

2.3.6 RAID Rebuild

When a disk in a RAID level 5, 6 or 10 array fails, something called a RAID rebuild has to be done. A RAID rebuild is, in practice, a data reconstruction. The failed disk is replaced, while the data (and parity files, if those exist) on the disk is copied over to another disk in the array. When the failed disk has been removed and a new one has been inserted into the array, the data is copied over to the new disk. The rebuild is important, since it is affected by multiple factors - for example, higher capacity drives, the bandwidth of the disks, or in the case of multiple disks failing at the same time (a mirrored image in a RAID 10 array for example), or when there are parity calculations that have to be conducted - all of these increase the rebuild time significantly. For a visual representation of how a RAID array might look like, see Figure 2.3.




Figure 2.2: A graphical representation of a RAID 10 configuration [2]



Figure 2.3: Storage servers utilizing RAID technology. [15].

However, as disk size increases, so does the chance for disk failure in a RAID

array. This, coupled with the fact that SSDs are on the rise, which do not work the same way as a conventional HDD and thus are susceptible to a lesser amount of disk failures, means that RAID technology might become obsolete in the coming years.

2.4 SAN

SAN stands for Storage Area Network. The main usage of a SAN is to store data, by increasing the accessibility of the storage devices within the network. It works by making all storage devices, which are often arranged in a RAID array, on the network available and ready to use by all servers on, for example, a LAN network. The size of a SAN can vary from a minimum of two devices, to many thousands.

There are several advantages to SANs. They offer great speed. For example, when the storage and the servers are connected using a network of high speed interconnected switches based on fibre channels, then it is possible to reach levels close to 5 Gbps [17]. Another advantage is that it is possible to achieve server-less backup of the data. This is achieved by letting disk storage devices copy their data onto so-called backup devices across the network, and this happens without any intervention from a server. This means that the data is stored exclusively on the SAN itself, and does not leak out into other networks, and thus does not consume bandwidth from other processes on the network. Finally, another advantage is that a SAN is very scalable, since it is possible to simply add storage devices when the need arises.

2.5 NAS

NAS stands for Network Attached Storage. Being different from the SAN described above, a NAS is often just a small computer, that has some differences from a regular computer, and which houses an array of storage devices in RAID mode. The difference can be summarized as follows [44]; "A NAS is a single storage device that operates on data files, while a SAN is a local network of multiple devices". A NAS only provides file-based storage, which means that you cannot run conventional software on a NAS. Furthermore, a NAS seldom uses keyboard and mouse, and often not even a display, and are controlled externally, through the network. Finally, a NAS does not have to have an operating system.

A NAS operates over different network protocols. For instance, HTTP, FTP and NFS are often used, and it is rare for a NAS to limit the user to one of these protocols. Indeed, many are often used in conjunction.

Some advantages of a NAS include the ability to be a helping hand for smaller systems, such as e-mail servers, by providing storage. Being a smaller system than, for example, a SAN, it is also possible for smaller consumers, such as private persons, to utilize a NAS when their storage requires it, and this usage has seen some recent rise as the prices for NAS devices has gone down. Finally, as stated, being a smaller device, the physical space occupied by a NAS is significantly smaller than the alternatives.

2.6 DAS

DAS stands for Direct-Attached Storage. A DAS is a digital storage that is, as the name suggests, directly attached to the computer on which it is going to be used. This is in stark contrast to both the SAN and the NAS mentioned above, where the storage is accessed through a network. In fact, most storage mediums utilized by personal computers fall under the DAS-category, such as the aforementioned hard disk drives and solid state drives, but also optic disk drives, and the DAS-category was created especially to separate these sort of storage devices from the ones mentioned earlier.

2.7 Magnetic tape

Magnetic tape is one of the oldest digital storage media, and still sees widespread usage today for archiving purposes. The main reason is the low cost compared to other storage mediums, as well as the very low power consumption.

Due to the nature of magnetic tape, to access data the tape needs to be moved to the correct position, in contrast to hard drives which uses random access. There is also no file index on the actual tape, as the index is often stored in an external database. This can prove to be a problem, since the data needs to be located before it can be read. To be able to accomplish this, specialized software has to be used in conjunction with a tape drive to be able to read and write data to the tape. Without this software, the data might be considered to be lost, since it cannot be located. This is also a major concern regarding redundancy, since if this database is corrupted, or the hardware on which the database is running goes down, the data on the tape can be considered lost. This contributes to tape being less portable than for example hard drives, where data can be located not only quickly, but without the help of an proprietary index. This is also in contrast to video tapes, where video can be replayed back with a video tape recorder, and does not need an index [52].

Most commercial tape setups today are based on the LTO-standard (Linear Tape Open).

2.7.1 The LTO standard

Today most tapes in use are, as mentioned, so called LTO-tapes, which is an open standard for magnetic tape. New LTO generations are released every 18-24 months, where each generation brings faster writing and read speeds, as well as larger storage capacities. At the time of this thesis, the current generation is LTO-8, which has the following specifications for a cartridge:

- 12 TB Uncompressed storage (30TB compressed)
- 360 MB/s reading and writing (uncompressed)
- 900 MB/s reading and writing (compressed)

Compared to mechanical hard drives, magnetic tape has very good reading and writing speeds, both compressed and uncompressed. The biggest drawback is

however that the tape needs to spin to the desired physical location on the tape before writing or reading can begin. This can take up to tens of seconds [52]. There is also a concern regarding how tape will perform when playing accessing multiple files at the same time. Since files are simply appended sequentially, and no random access exists, the tape would probably not handle this use case very well. The tape drive would have to read data from many different locations on the cartridge at the same time, which would result in the tape being spun back and forth a large number of times. This would most likely decrease the tape lifetime as well as the performance.




Figure 2.4: Storage capacity of LTO and other tape technologies. Shows LTO up to LTO-7. Note that LTO-8 is now available in market [19].

Compression in LTO-tapes

All LTO formats include hardware compression of data. This is indicated by the specifications on the tape itself, where two different storage sizes are mentioned, compressed and uncompressed. See Figure 2.4 for an image showing the storage capacities of the different tape formats.

The compressed size can be a bit misleading, since it often assumes a 2:1 ration, which can be hard to achieve. In fact, the compression might lead to larger file sizes for some input due to the pigeonhole principle. If 2^N bits are compressed into a sequence of K bits, where $K < N$, then only 2^K sequences may be produced. This will lead to different sequences compressing to the same sequence, and thus

they can not be uncompressed.

This is the case with pre-processed multimedia content such as MPEG, MP4 and JPEG etc. When using hardware compression for these file types the size will often increase to a larger one for the "compressed" file than the uncompressed one. Thus the compressed file size indicator on the tape can not be used for multimedia content, and the uncompressed size should be used instead.

2.7.2 Tape libraries

The nature of tapes poses yet another problem. To be able to read and write data to tape, the tape needs to be physically inserted into the tape drive. This has led to devices called tape libraries to be developed, that can load tapes automatically without human intervention. It does this by stamping tapes with a barcode that the system can recognize. It then uses these barcodes to find the correct drive when the data needs to be read. These tape libraries are very expensive however, and are typically used by larger business that house a lot of data. See image 2.5 for a typical tape library system.



Figure 2.5: A StorageTek tape library system [18].

2.7.3 LTFS

One of the biggest drawbacks with tape is on how data is being read and written. Since the tape does not provide random access in the same way as an HDD, the tape has to be physically spun inside of the drive until the starting point of the data is found. To read the file, the tape drive then has to read the tape from the starting location of the file sequentially until the end. This results in a need for specialized software to be able to read and write data to the tape, and hence the LTFS (Linear Tape File Format) was created — to make magnetic tapes portable and independent of a proprietary index. The main aim for the project was to make tape as easy to use as hard drives, meaning that files should be viewable to the user without the use of an external index, so that the user can interact with the files on the tape via a file explorer. Directories can be opened and files can be used as if they were stored on a hard drive, with the exception of the storage medium behaving like a tape drive, meaning that there exists no random access [52].

It is however important to note that even though the tape acts as a random access storage medium, it is still tape, and therefore still comes with the limitations of tape. Using tape as a hard drive can prove to be a painful user experience for interactive use, where loading times to access a file can be perceived as very long while the tape is spinning to different file locations.

Nonetheless, LTFS brings simplicity to the tape format, where tape can easily be used without any proprietary backup software. It still comes with the limitations of the physical format, with long seek times.

2.8 MAID

MAID (Massive Array of Idle Disks) is a storage technology that aims to be an alternative between RAID drives and tape libraries. The motivation behind this is to keep the benefits of using magnetic tape, while reducing costs, power consumption and heat generation. Other benefits include faster data retrieval times compared to tape libraries [30]. The main idea behind MAID is that not all disks are active at the same time. Disk are turned off until the data they hold is requested, and they are spun up again. This spinning up process takes time, and therefore frequently accessed data is moved over to cache disks, to speed up access time.

The main assumption behind MAID is that less than 5% of the data in the system gets requested daily. Therefore MAID is a good choice for archiving system, where it has seen some success[42].

However, MAID never saw much greater success, and it seems to be hard to find commercial systems for sale. The power savings that were promised were exaggerated, and the savings were not enough to justify the MAID system. It seems like the MAID format is making its way back into data centres recently, with systems such as Microsoft Pelican, which is set to deploy in Microsoft's Azure data centers. The Microsoft Pelican houses up to 11.5 petabytes, and keeps only 8% of its drives active at any given time, reducing power consumption and reducing the need for extensive cooling solutions, where the main target group is for systems

too small for tape libraries [23]. Microsoft's recent interest in MAID system might be an indication of the MAID architecture as a whole becoming popular once again for data centres when instant access is not essential and tape libraries are too complex to use.

MAID systems might be useful for data that is never or very rarely accessed, or stored in a specific way, such as video surveillance, which is very sequential in nature, both in accessing and storing. The power savings might be even greater since video surveillance data is very rarely accessed, which could lead to even less drives being powered on in a MAID system.

2.9 Cloud Storage

A relatively new form of storage that has become prominent in the last few years is the so-called cloud storage. Cloud storage is a storage model in which data is stored on remote servers, that are accessed through a network (most commonly through the internet), often referred to as the "cloud". The cloud storage is, in the majority of cases, not maintained by the users themselves, but are paid for as a service. The storage is then accessed through a network-based API, which the users can access at any time.

Users often pay for the service by the amount of data that they store. It is common for services to charge a price per gigabyte. Sometimes there are other expenditures as well. For example, some services make users pay for the retrieval of data, and users sometimes have the option to pay extra to be able to retrieve the data at a faster rate than what is normally available [5]. The users are often charged by the *amount* of data they wish to retrieve as well, not only for the retrieval itself. Retrieval times are another thing to consider with cloud storage. Since the storage devices are not physically present at the location where data might want to be retrieved, longer retrieval times become a reality and retrieval times can vary from between minutes to hours. For more concrete numbers, see Sections 5.2.1 and 5.2.2.

Security and safety is often another concern when it comes to cloud storage. A study conducted by Clutch [49] reported that "forty-seven percent (47%) of respondents are only "slightly" or "not at all" comfortable with storing personal information on Apple's iCloud." It is speculated that the "[...] hesitation may be a result of [the users] limited knowledge of iCloud." But then, it is reported that a majority of the users who *are* familiar with the functions of iCloud are still only "somewhat comfortable" to "not at all comfortable" with using the service [49]. It should be noted that in the recent years, there has been breaches into the personal accounts of some celebrities, who have had their private photos leaked, and it is possible that the media coverage of these leaks has had an effect on the public opinion of cloud storage for personal use, even though iCloud boasts multiple security features, such as encryption and advanced password requirements [48].

When it comes to industrial usage of the cloud and its abilities, it remains more widespread. A study performed by scientists at Luleå University of Technology in Sweden showed that, out of 110 qualified answers received from the study, coming from people regarded as knowledgeable within the industry, 99% of all businesses

either are currently using or are planning to make use of cloud computing. Cloud computing and "the cloud" is here referred to as "[...] the Internet, and in practice it refers to the data center hardware and software that support the clients' needs, often in the form of data storage and remotely hosted applications" [25].

2.10 Summary

In this chapter, several different storage technologies such as hard drives and magnetic tapes were presented along with their properties, advantages and disadvantages. The standard HDD was introduced and briefly described, along with the newer SSD, and their specifics and differences were noted. Then RAID technology was introduced, along with some of the most common storage network technologies, such as SAN and NAS. RAID, in combination with some of these network technologies, can greatly increase the performance, for example, in terms of speed and redundancy, of storage networks, if implemented correctly.

Tape storage is slowly making a comeback due to the large amounts of data it can potentially store. Magnetic tape offers very low costs per terabyte, as well as very low power consumption. At the same time it can be a very cumbersome technology to work with compared to other more widely used storage technologies, such as hard drives and NAS systems and requires special implementation as well as manual handling.

The concept of spinning down hard drives to conserve power, in some instances referred to as MAID, was presented to give the reader some acquaintance with this kind of technology. MAID was supposed to be a storage with properties somewhere in between traditional hard drive storage and slower but cheaper tape storage. MAID never really saw widespread application, mainly due to few producers and lower energy saving than what was predicted. However, MAID-like technologies are starting to see a return to form with large corporations such as Microsoft starting to take interest in similar technologies.

And finally, cloud storage was introduced and discussed. As with the other storage technologies, there are several things to consider with cloud storage, both advantages and disadvantages. Cloud storage comes in many different shapes, some being constructed for use by private customers, and others for commercial use for bigger clients. The price for cloud storage is often billed monthly, and with a cost per piece of data retrieved and stored, as opposed to the other storage options where you pay once for your storage medium and then use what is available on it. There is also often concerns with the security of cloud storage, and even though there seems to be great effort put into this concern by the industry, public opinion is still not entirely convinced.

Analysis of properties related to storage

In this chapter, we analyze different properties related to common storage technologies, and with that different metrics which could be argued are of great importance when selecting a technology to work with, such as price, performance and security, and why we chose these metrics. These evaluations served as a basis for evaluating the products and services in later chapters.

3.1 Cost of storage

Cost is arguably an important factor when making any choice related to technology. How high of a cost is acceptable with regards to the sizes of the organizations which are utilizing said technology? It has to be noted that our use case states that these organizations will be required by law to house one of these storage systems. The system is not, therefore, essential for the business that they are performing, and so the price has to reflect this.

When the prices were evaluated, we looked at the problem from a perspective of price per terabyte (\$/TB). This perspective was chosen since it gave us a simple metric to compare the cost across many different storage technologies and price models.

As we were comparing the products, it was noted that some products came with a one-time cost of purchasing the equipment, while others employed a model based on subscription fees. For example, most of the physical storage products simply had no subscription model, while most of the cloud storage services had. The subscription based models used by cloud storage services were based on the amount of data that was stored, and on the amount of times that the user chose to retrieve data from the cloud. Similarly, the physical storage devices of course also has a max amount of data storage. The requirements of the system (see Section 1.4.2) gave us a rough estimate of about 500TB per year of storage.

The physical storage products require the purchase of a product, which is often a form of drive, and also the purchase of storage devices which are inserted into the drive, onto which data is stored. The storage devices are almost always in the form of HDDs or SSDs, except in the case of tape drives, when the storage devices are in the form magnetic tape. When these prices are calculated, care has to be taken to both the price of the drives themselves, and also to the cost of the storage units. Different physical storage products can use different amount

of storage disks as well; for example, one device might be able to house 12 disks while another can house 20 disks.

3.2 Power consumption

Power consumption is a metric which is closely related to cost. However, it is also related to sustainability. When talking about power consumption, or energy consumption, the distinction is made between direct energy and indirect energy [54]. *Direct* energy is what is meant by sources of energy which is being depleted by the user, such as gasoline or electricity, while *indirect* energy is often described as energy which is embodied inside consumer goods. Using these terms, the energy that is being consumed by cloud storage services would be classified as indirect energy, since it is not being consumed directly by the user, but rather by the servers at the site of the service provider. Energy being consumed by devices in the form of drives and their storage, on the other hand, would be classified as direct energy.

Measuring the direct energy that is consumed in terms of a pure economic context is trivial, all that is needed is simply to know how many Watts a device is consuming. For example, one of the storage devices that was investigated consumed 511W for typical usage. Since the devices are powered on all day (with the exception of some of the magnetic tape drives), every day, to be able to store the streamed data, we can convert Watt to kWh (kiloWatt-hours) with a power-on time of 24 hours. The device in the example then consumed, during one year, 4476.4 kWh. Another device we investigated consumed 525.6 kWh during typical usage for one year. As can be seen, the difference is quite substantial.

It is harder to put a price on the indirect energy that a service consumes, since the price is often included in the service. A cloud storage service provider has a certain cost that is being paid by the users of the service.

Measuring the power consumption in terms of its effect on the climate is harder still. It can here be relevant to compare the energy consumption of a physical storage device, like the ones mentioned above with, for example, a refrigerator, since both devices are powered on all day, every day. A standard refrigerator (with a volume of around 400-500 liters) consumes about 495 kWh per year [24]. In the case of the storage device which consumed 511W and 4476.4 kWh per year, it can therefore be argued whether or not it is really justified to keep a device which uses as much energy as about 9 refrigerators, just to store data which in the majority of the cases can be regarded as cold data (see Section 4.3), meaning that it is not often, if ever, accessed.

Furthermore, the best thing an organization can do is to make sure that the source of the energy used for their storage devices (i.e. their direct energy usage) is coming from renewable energy sources, such as solar power, wind power and hydro power. If possible, the same should be applied to indirect energy usage, and some service providers, like Amazon, claim to be working towards the goal of using primarily renewable energy sources [4].

It is fairly easy to transform the power consumption to a cost. By looking at the average price of a kWh [3], and then just multiplying this price with the

number of kWhs for a system, we can see that the cost of the power consumption for the system with the lowest number of kWhs per year, was $\$0.10 * 210.24 \text{ kWh} = \21.024 per year, which was for the magnetic tape drives. It should be noted that this measurement, 24W which equals 210.24 kWh per year, is in the case of the tape drives being powered on all day, every day, which they probably won't have to be, and so the cost could probably go down even further.

Another thing to consider when talking about systems using a standard RAID solution is that these will require some form of cooling. Almost all of the energy consumed by the unit will be converted into thermal energy. This thermal energy will have to be moved away from the storage unit - mostly by fans - and into the server room itself. The temperature will increase in the server room, to that of the storage device itself if no ventilation is provided.

The equation $Q = \Delta Tmc$ describes the relation between the heat energy (Q), the change in temperature that can be observed (ΔT), the mass of the object (m) and the specific heat of the object (c). Since the power consumption of the storage system will almost equal the Q in the equation above, it becomes relatively easy to calculate the temperature increase. It therefore becomes important to provide adequate ventilation. The maximum recommended temperature for server rooms is somewhere around 28°C. Hotter than this might start to affect the drive's lifetime.

3.3 Drive failure & lifetime

Drive failure is a problem for any organization storing large amounts of data digitally. There are many organizations, and private persons, storing data digitally; according to Google researchers, about 90% of all new information produced in the world in 2007 was being stored on magnetic media (HDDs) [53]. It is reasonable to assume that that number holds largely true today as well.

3.3.1 HDD & SSD

Drive failure typically refers to the sudden inability for the drive to function properly. Drive failure can happen due to a number of reasons. It is commonly thought that a common reason for drive failure is temperature; too high temperature to be specific. Indeed, as stated in the previously mentioned study [53]: "Temperature is often quoted as the most important environmental factor affecting disk drive reliability". However, it was shown in the same study [53], as well as by Backblaze in 2014 [28], that there appears to be no correlation between drive temperature and failure. Nonetheless, it is worth to note that a similar study was conducted by the University of Virginia on a Microsoft data center, in 2013, which *did* find a correlation between too high drive temperature and failure [57], so there appears to be some need for further research on the subject, since the current research has produced conflicting results.

Another common reason for drive failure, in the case of magnetic HDDs, is vibrations [59]. Vibrations disturb the reading and writing of data on the magnetic storage by the rotating disks, and can, if severe enough, cause the disks to shift slightly and damage the storage units, causing the drive to fail. As previously mentioned, this is not a problem if the drive is an SSD, since it does not have any

moving parts and therefore is not susceptible to the same conditions. In relation to the vibration issue, something that can occur to all types of drives and has a potentially damaging effect is of course excessive physical force, such as when the drive is dropped. Again, when it comes to magnetic HDDs, they are also susceptible to magnetic field damage. Other electronic devices may emit magnetic fields with strength enough to damage the magnetic drives of the HDD.

It is also not uncommon for drives to fail due to the power going down. Drives which store data continuously require uninterrupted power supply (UPS), and when the UPS is interrupted, drive failure becomes more common, since the drive might be in the middle of an operation.

When it comes to lifetime expectancy's, for the HDDs and SSDs, more information can be found in Sections 2.1 and 2.2, respectively. It is worth noting, though, that certain solutions and technologies can extend the lifetime of the disks used. Some systems utilize a special technology of spinning down idle disks, thus keeping them inactive for a majority of the time. According to one company utilizing this technology in their products, it can extend the lifetime of disks used by up to ten times. This would mean that the lifetime would be, if not at least the same, then longer than that of the magnetic tape drives.

3.3.2 Magnetic tape

The three contributing factors to tape failure are loss of magnetism, degradation of coating materials, and dimensional instability [37]. These are not the only factors that are mentioned when it comes to tape stability - temperature is also considered to be an important factor for the stability of the tape cartridge. Temperatures over 23°C can increase the tape tightness, leading to permanent data loss in the tape. Humidities over 70% can lead to expansion of the tape, as the tape absorbs the moisture from the air. High humidity can also lead to mold growing on the tape, since the mold can feed of the binder polymer, and lead to data loss on the tape [11]. Magnetic tape is, as mentioned, also prone to degradation due to magnetic interference, which can destroy the contents of the tape [37].

Magnetic tape has a long lifetime, but as can be seen, the requirements on the storage environment are higher compared to other storage media. The tapes are also rated for how many passes can be performed on the tape before it fails. These reach high numbers and should not be a cause for concern. One tape can handle approximately 260 full backups [47]. The lifetime of tape is also very long when compared to other storage mediums; it is around 15-30 years, which is much longer than what is needed for our use case.

3.3.3 Cloud storage

Cloud storage, which is really just a combination of HDDs, SSDs and other forms of storage devices connected to one or more networks like the ones mentioned in Chapter 2, at another site, share the same failures for the drives as the ones mentioned above. But the difference is that cloud storage service providers often boast large backup units, to guarantee redundancy for their clients.

In the case of cloud storage, how long can the servers of the cloud storage

service providers be expected to last? Of course, all the service providers will aim to keep their clouds up for as long as possible, but does this mean indefinitely? The question is hard to answer, since technology itself moves so fast, that it is hard to say if cloud storage will even be around in a few decades. This, coupled with the fact that there aren't really any good examples of cloud services being around for more than a decade (one of the oldest services, Dropbox, launched in 2007), makes it hard to say if they will exist in the far future, at least in the same sense that they do today.

3.3.4 Cost aspect

Furthermore, there is a cost aspect to this analysis as well. If a higher redundancy is wanted, for the physical systems, then the only thing which can be done is to simply acquire several drives. For example, with the tape systems, if the user wants to have higher redundancy, several drives need to run at the same time, to spread out the data, which effectively at least doubles the cost. The same can be said for other physical systems as well. In addition to this, if a drive failure does occur, then this is sure to cause some financial setback as data recovery can be quite costly. If permanent data loss is the case, then in the worst-case scenario, the result can be bankruptcy – see Section 1.3.1.

3.3.5 Geo-redundancy

A case can also be made for what is commonly referred to as geo-redundancy [38]. What is meant by geo-redundancy is that data is placed at two separate physical locations. This is to protect the data in the case of catastrophic events such as earthquakes, fires and tornadoes; if one server containing data goes down, another one is still running with the same data, and so no loss of data occurs. It is reasonable to assume that cloud storage service providers employ some form of geo-redundancy at their sites. In the case of the smaller, physical storage systems that have been discussed in this report, it would on the other hand be unreasonable to work towards geo-redundancy - in this case, it would mean that the user would have to set up at least two storage-stations at different locations, presumably one at their own site and another somewhere else, and that they would have to maintain both, which result in extra costs both in time spent and in maintenance equipment, not to mention that at least two sets of hardware have to be purchased.

3.3.6 Summary

The best protection against immediate failure, then, comes from cloud storage. This is mainly due to the redundancy offered by the service providers, which ensures that users can always reclaim their data, should they want to do so. If pure longevity is considered, the best storage medium for video data is some kind of film reel [40], like 8mm, however, this is highly impractical and not feasible at all for the purposes that we are discussing. Nonetheless, considering our use cases, in which the data needs to be stored up to 1 year, we are not really interested in data having to be storable for more than that period. If a paradigm shift is approaching with regards to storage mediums, the relatively short time (even though we are

talking about *long* term storage) means that such a shift could most likely be handled safely, if cloud storage is utilized.

3.4 Scalability

The scalability of a system refers to its capacity to change in size, but very often with storage systems, what is meant is the capacity for it to grow, since there is often a need to store more data. With the systems we have investigated, there is a difference with regards to *how* they scale. Using magnetic tape storage, it is possible to increase its storage capabilities by simply adding more storage tapes for use in the tape drive(s). This is called *vertical scaling*, or scaling up, and what it means is that more storage is added to the existing hardware infrastructure.

Using a NAS system it is possible to increase the storage capabilities of the system by adding more hardware in the form of storage systems. This is called *horizontal scaling*, or scaling out. So with horizontal scaling, more hardware in the form of individual storage units has to be added [21].

Cloud storage doesn't really fit into any of these models, since no hardware is involved, for the user. For the cloud storage service providers, again, they are secretive of their work with regards to how the scaling is performed, but most likely scaling is done both vertically and horizontally since the services need to accommodate a lot of data.

It is worth to note, however, that there is another difference between the vertical scaling of the tape drives and the horizontal scaling of the NAS systems. Today, the size of, a common size of storage disks is 14TB. Seagate is working towards releasing a 20TB HDD in 2020 [16], and according to one source [22], models of 40TB and above are to be expected sometime 2023. Seagate themselves even state that capacities of 100TB will be possible in the future.

If these predictions come true, then it will most likely be possible to use these larger HDDs in a NAS system. If for example 20 TB disks are utilized in a system which houses 15 disks with sizes of 14 TB for storage purposes, then the total capacity would increase by $15 * (20 - 14) = 90$ TB, from 210 TB to 300 TB. Even larger HDDs would give an even larger increase in capacity, with no additional hardware having to be acquired, and the larger the storage capacity of the NAS system housing the disks, the larger the increase in total storage.

Magnetic tape technology works differently; when new standards of LTO are introduced to market, like LTO-9 as described earlier (see Section 2.7), new hardware in the form of tape drives will have to be acquired. This means that whenever the storage medium, i.e. the tapes themselves, grow in capacity, then to meet that new capacity more money has to be spent. This is in stark contrast to the above mentioned technology, where the same storage systems can house storage devices of the same technology but with higher capacity.

3.5 Performance

Performance for a storage system can mean different things. One common metric is *throughput*, which means how fast the storage system can deliver data. Another

is *retrieval time*, or *response time*. This is how fast the system can access data. This can be measured from certain perspectives, like the user perspective, which is the most common. Of course, another performance metric which cannot be ignored when discussing storage systems is *capacity*, with regards to how much data the system can store [29].

Throughput in storage systems is usually measured in Bytes per second (Bps), or bits per second (bps). Some technologies, like RAID, can increase the throughput of a system. Utilizing RAID level 0 - distributing the contents of one file across all disks in the network - can actually increase the throughput of all operations, both read and write.

It is hard to say what the throughput is when there is no information about the hardware utilized, as is the case in most cloud storage solutions (as mentioned, Amazon, for example, is very secretive about their work), and so, we have to rely on measuring the retrieval times of those systems. We can safely assume that most cloud storage providers utilize some form of RAID - or similar - structures. The throughput for the magnetic tape standard that were investigated in this report can be found in Section 2.7.1.

When it comes to retrieval times, these are more important from a user perspective, since they measure how long the users actually have to wait before getting their data. In most circumstances, one wants the retrieval time to be as short as possible, usually in the range of seconds. However, in some cases such as for data that is very rarely accessed, retrieval times of several minutes or even hours are deemed perfectly acceptable. In this case, retrieval times of up to one or more hours are acceptable.

As previously mentioned, the retrieval times for cloud storage systems can vary, between minutes and hours. Retrieval times for physical, local systems are faster and are practically instantaneous in the same context - the retrieval time may go up to the tens of seconds, at most, with tape being the exception, where there might be a need to find the correct cartridge. This is discussed in greater detail in Chapter 5.

Capacity then can of course be regarded as perhaps the most important metric with regards to performance for a storage system. A higher capacity means that the user has to invest less time managing the system, since the storage capacity is then less likely to be filled. Indeed, a high enough capacity coupled with older data being sequentially overwritten by newer data means that the user may never have to interfere with the system at all, lest of course it is required by law officials should the need arise. The capacity needed for a long term storage system for video surveillance for the system we investigated ranges in the 100s of TBs (terabytes) of data.

Cloud storage has potentially infinite storage capacity, since the user can simply purchase more capacity as the need arises, but as mentioned, then the cost increases. Physical systems have a set amount of capacity, and so it is more important to make the best use of the capacity in this regard. The manufacturers which we investigated had products with capacity ranging from 30TB up to 630TB.

It is not easy to say which technology is the most suitable with regards to performance, since different technologies perform better at certain aspects. Cloud storage has the capacity factor, physical units have better retrieval times and so

on.

3.6 Accessibility

With the accessibility of a system, what is meant is ease of implementation or installation, ease of use and ease of maintenance. All of these aspects are important from a technological perspective, since they form a measurement of how much time is needed for the system to run and work as intended, and as such, they can be regarded as a way to further measure the cost of a system. Below, they will be discussed with regards to the storage systems that were investigated in this thesis.

Ease of implementation refers to how easy the system is to install and set up the first time it is introduced. For the users of the finished system, this is not interesting, but from a developer perspective it is, since a difficult installation might take longer time, and more time spent on each installation equals a higher cost. Ease of use refers to how easy the system is to work with on a day to day basis. This perspective is interesting for both the users and the developers - each group wants a system that is as easy to work with as possible, of course. In reality, a system such as this should ideally not see much use from a user perspective on a day to day basis, since its only purpose is to store and archive video data. Ease of maintenance refers to how easy it is to maintain the system, and eventually repair it if the need arises. This is another factor that is more of an issue with regards to developers, since they will be the ones performing the maintenance.

Special note has to be taken here with regards to the cloud storage service providers and their systems that we looked at. Since cloud storage systems lack any sort of physical infrastructure from a customer perspective (user perspective), the only two accessibility aspects of note are ease of implementation and ease of use. Most cloud storage systems provide the users with an API that lets them connect to the cloud servers. In general, installations of this type are not very difficult to perform and thus, are not very costly in terms of time and money spent. Maintenance falls entirely upon the cloud storage service providers.

Physical storage systems, including magnetic tape, are subject to more problems due to hardware issues that may arise - as previously mentioned, disks can fail in many ways, for example. If a storage device such as this fails, a service technician may have to be dispatched, and so the maintenance cost of such a system will be of greater interest. Therefore, ease of maintenance is of course an important factor with regards to physical storage systems.

With all these aspects in mind, the most accessible system is deemed to be a cloud storage system. Several issues, such as maintenance of the system, physical accommodation of the system, and issues related to scalability, falls entirely upon the cloud storage service providers.

3.7 Security

The security of the stored data is another aspect which is good to take into consideration, and again, distinction has to be made between physical stored data and cloud stored data. With regards to physically stored data, the data is to be

archived on several disks that are located at the customer. In the case of magnetic tape, these tape drives are not connected to any network and the tapes themselves are physically located on-site, and thus this is in a sense the most secure form of storage. This is of course, not regarding factors that are outside of the control of what can be provided by this sort system, like break-ins and disasters. Other physical systems utilizing some sort of RAID structure in combination with one of the network structures also discussed earlier are more vulnerable since they are susceptible to attacks. SAN networks, for example, can be vulnerable to attacks like man-in-the-middle and session hijacking.

Cloud storage has already been briefly discussed with regards to security (see Section 2.9). Cloud storage is secure in a number of ways. Redundancy, as has been mentioned, provides safety against disk failure. Furthermore, the data on a cloud network is located on a server in a remote physical location in a data center, and these are often very secure and protected with regards to unwanted personnel. This means that if anything happens to the local files of the user (for example, if they are corrupted by a virus or any other form of malicious software), they may be regarded as lost - but this is not the case with data stored on servers hosted by a cloud storage service provider.

The fact remains, however, that with cloud storage, you never really 'own' the data in the same way as you do when you own the physical disks and drives which houses the data, since the data is present at the storage site of the cloud service provider. It is still yours in the legal sense, but you cannot pick it up and do with it as you wish without first extracting the data from the storage site and putting it in some local storage. In a sense, this is what the user is paying for; the risk that comes with storing data can, by utilizing cloud storage, with relative ease be moved to a service provider that has certain guarantees with regards to security (and of course, as mentioned, accessibility), in exchange for a price.

3.8 Summary of evaluations

In this chapter we've discussed different metrics which were used to evaluate the storage systems that were investigated. Specifically, in the cases where it has been applicable, the cost has been tied to the metric in case. \$/TB has been a rigid metric since it allows for a direct comparison for the cost of storage between the different storage systems and payment plans. Power consumption as a metric has also been discussed, and a distinction has been made between direct and indirect energy.

Drive failure & lifetime expectancies were also discussed. Some common reasons for drive failure were exemplified, and the lifetime of different storage technologies were mentioned. Redundancy was discussed as a protection against drive failure, and geo-redundancy was also discussed.

Scalability was another metric which was discussed, and the scalability of the different systems were compared with regards to how they scale - horizontally or vertically. It was further stated that there is a difference with regards to the storage mediums of the different systems, and how that difference effects the cost of the scalability.

Performance as a metric was discussed with regards to throughput, retrieval time and capacity. The performance of a system is closely related to certain costs; if a system performs worse with regards to for example retrieval times, this naturally means that less retrievals can be performed during a fixed time interval - and this of course applies to all forms of executions in the system. The more time is spent waiting for results, the less actions can be performed.

Previous and relevant work for data archiving and data centres

In this chapter, the methods used by corporations that handle big data will be presented. How do they handle data that might not be accessed frequently, but still has to be saved for compliance, or if their users request it?

The impact that the data storage sector has on the environment will be discussed. Relevant work into creating energy efficient and cheap storage solutions will be presented.

4.1 Big data

When it comes to storing data for long periods of time, the most common approach today is to store the data in large data centers with thousands, if not tens of thousands, of storage devices. These storage devices are often in the form of the above mentioned HDDs or SSDs, but sometimes they are also in the form of magnetic tapes. Furthermore, to improve the storage capabilities further, the storage devices are oftentimes arranged in RAID or RAID-like structures, and a SAN or a NAS is applied to the network.

With the rise of increasing storage demands and requirements, the companies offering cloud storage are developing new and intriguing technology to cope with these demands. These new technologies promise lower costs per terabyte, as well as lower operating costs in the form of lower power consumption.

4.2 Environmental impact of data centers

It is important to shed light on the impact that data has on the environment, a topic that many are not too familiar with.

Estimating total power consumption of the world's data centers is not an easy task to do, though there have been several attempts to do so. According to the US Department of Energy, data centers in the US amounted to 1.8% of the country's total electricity consumption in 2014, and the current projection is that the power consumption of data centers in the US will reach 73 billion kWh in 2020 [58]. However, as more and more people around the world gets access to the internet,

storage requirements will likely increase even more, as well as energy consumption for data centers.

As Peter Gross, who has helped design several data centers, said: "A single data center can consume more power than a medium sized town" [36]. For example, Facebook's data center in Luleå, Sweden, consumes as much power as that of the municipality of Västerås in Sweden [32].

It is therefore important to increase the energy efficiency, and thereby decrease the amount of energy used in data centers, as much as possible. There is also an incentive for data providers to decrease the power consumption of their data centers, as the electricity bill not only decreases with the decreased energy demand from the storage itself, but also from the cooling. With the reduced heat, the storage density can also be increased, meaning even lower costs.

4.3 The tiered storage model

With the rise in storage demands as well as retention requirements, there is an increasing need to divide up data into different tiers. The basis of the tiered storage system is that the most frequently used - and the most vital - data is put on the fastest storage mediums possible. Other data is put onto other slower but cheaper storage mediums, since it is not accessed frequently. This gives rise to a tiered storage model that is frequently used to model storage systems. In short it can be summarized as following:

- Mission critical data
- Hot data
- Cold data

These different tiers have different requirements for retrieval times as well as storage space. Usually, mission critical and hot data is put on hard drives or similar storage systems, while cold data is put on magnetic tape because of its low price [34]. However, as mentioned in previous chapters, magnetic tape can prove to be cumbersome to work with compared to a disk based storage system. Therefore a number of other interesting investigations has been made into making disk based systems more energy efficient [27].

The cold data tier is of especial interest since the access requirements are significantly lower than that of hot or mission critical data. A case can be made for classifying surveillance video as cold data, since it is very rarely accessed.

4.4 Pergamum

Pergamum is a storage system created by researchers at the University of California. Since so much data is being created, long term storage must keep up with the growing storage requirements. Also, new laws put increasing pressure on actors to save their data to comply under data retention legislation.

They claim that many storage options today favor performance over energy efficiency, which is not needed for an archival system. Previous work into MAID

systems that was previously described, does not fully utilize the power efficiency that is possible from an idle disk system. The Pergamum managed to achieve lower power consumption than the Copan MAID system, one of the few MAID systems that were commercially available.

They managed to develop Pergamum in such a way that the power consumption was comparable to that of tape libraries, while achieving higher reliability and faster random access [60].

4.5 Microsoft Pelican

As previously mentioned, Microsoft has constructed a storage system that they refer to as the Pelican. Pelican is an attempt to increase storage density, decrease costs, and decrease power requirements for big data systems, where the data is not frequently accessed. Like the MAID systems mentioned in Section 2.8, the Pelican operates on a basis of spinning down disks when they are not being used. The rack has space for 1152 disk, but only enough power supply to keep 144 disks spinning at a given time. The biggest difference is the hard constraint of 144 active disks at a time, in contrast to MAID systems where there was no hard limit on the number of active disks. The other disks are in idle mode, where they are still powered, but the disks are not spinning. With these constraints, the storage density can be increased, resulting in a cheaper product in both acquisition as well as operational costs [33].

4.6 Facebook Open Vault

Facebook has also started to look into cold storage solutions, and have done so with their Open Vault project. Facebook is an interesting actor in this space because of the nature of the data that they store. When new media is posted to Facebook, it is accessed frequently. As the media ages, it gets accessed less frequently, with the oldest photos being accessed very rarely. In fact, 8% of photos gained the most traffic [43]. It would be a benefit to be able to move this cold data onto something cheaper - but slower - to be able to save power.

Facebook constructed a new kind of server warehouse, where servers were powered up as needed. This reduced power requirements, and they could serve the warehouse with one sixth of the power compared to a traditional warehouse. They also introduced constraints on that only one drive per tray can be active at any given time. This reduction in power consumption meant less heat, and thus the number of fans could be reduced from six to four, further reducing the power consumption [26].

Facebook has also made experiments in Blu Ray archival storage, with a setup that very much resembles that of a tape library, with thousands of Blu Ray discs. This system is not for primary use though, as it takes time to retrieve data from the Blu Ray discs, but serves Facebook's needs in the colder tiers as it is very power efficient, and costs half of what regular drive storage would [43].

4.7 Summary

In this chapter, previous work related to this thesis was presented. Many different organizations are starting to pursue alternative methods of storing data, compared to the more traditional methods of simple RAID storage.

Pergamum was a research product at University of California, trying to improve upon the previous MAID systems, which it did with success. It was shown that power consumption could be reduced while still achieving acceptable data retrieval times for archival storage. However, there is still no commercially available product built on this kind of technology that we have seen to this date.

Microsoft has also entered this field with their Pelican project. The Pelican has tried to improve upon the Pergamum design, while adding hard constraints on how many hard drives can be active at any given time. This led to increased storage density and lower power consumption. Microsoft is one of the few cloud providers that have publicly released information about the work that they are performing in this field, while Amazon and Google remain silent. It is reasonable to assume that both of them use some kind of similar system, or a hybrid system with tape libraries involved.

Facebook has introduced what they are doing in this field with their Open Vault Projects, since they have large amounts of cold data. They have done experiments similar to those that Microsoft has done with their Pelican system, and they have also experimented with Blu Ray technology, bringing down costs even more.

It is clear that the big cloud providers such as Google, Facebook, Microsoft, and Amazon, are interested in this kind of hybrid technology, where either tape is combined with hard drives, or some kind of spin down and scheduling technology is used. Since data is increasing by large amounts each year, and a big portion of it never - or very rarely - gets accessed, there is a need for long term storage solutions that is cheap and power efficient. It seems reasonable to assume that more technologies like this will be available in the future, maybe even for commercial use outside of big data usage. For the moment however, most of these technologies are only proof of concepts for further development in this field.

Product presentation and technical evaluation

To be able to solve the long term storage problem, a market analysis had to be conducted to gauge the different solutions available. Since this is a fairly new problem in the video surveillance space, the different solutions use different techniques. The main problem of finding a product is that they are often targeted to large scale data centers, and therefore tend to be expensive and do not target surveillance specifically. As mentioned, it is important to take into consideration that the surveillance data is not something that is relevant to the user's business, but is rather a legal requirement to be allowed to operate.

This chapter will present and evaluate the technical aspects of the systems. In addition to cost, which will be discussed in Chapter 6, these metrics offer knowledge about the most important parts of a system. This chapter will describe the specifications of the investigated systems, as well as the implementation that would take place in order to get them to work together with the VMS. The implementation specifics can be very different since the technologies themselves are very different.

As such, the metrics that will be evaluated in this chapter are: (a) performance and (b) integration with the VMS.

5.1 The systems

In this section, the systems chosen for evaluation will be presented.

5.1.1 Cloud solutions

Today, using cloud systems for storage is very popular. Services such as AWS Glacier and Microsoft Azure have a large number of users. Cloud is also used for video surveillance data in some cases. Therefore it was important to investigate if cloud storage could be used as a long term storage option.

Many of the cloud storage vendors have started to offer cloud archive solutions. These include AWS Glacier and Microsoft Azure with Archival Storage, and these two were the solutions that we investigated. They offer much cheaper data storage costs than their regular cloud storage counterparts, with the caveat being slower

retrieval times. In contrast to the regular counterparts, retrieval of data is not free, and a cost per gigabyte is not uncommon. The method of how the vendors store their data is not public information. Theories regarding AWS Glacier include tape storage as well as some MAID-like spin down technology, since the retrieval times are very long.

5.1.2 Veracity Coldstore

Veracity Coldstore is a NAS device that is specifically targeted towards video surveillance storage. The Veracity Coldstore 3U Pro holds 15 regular drives that each can store 14TB of data, when using today's largest hard drives. This in turns gives it a capacity of 196TB in total, where the last drive is kept for redundancy in case of drive failure. Veracity also has a larger model named the Veracity Coldstore Colossus, which can store up to 630TB (616TB effective) of video data on 45 hard drives.

In ways similar to MAID described in the storage technologies chapter, the Coldstore line of products spin down drives that are not active. In the Coldstore products, the technology utilized is called LAID (Linear Array of Independent Disks). In contrast to MAID systems where data might be accessed frequently, Veracity claims that only 3% of video surveillance recordings are ever played back, and thus the need for drives to always be powered on is very unnecessary. Axis has said that older recordings are very rarely played back, while newer recordings are played back very often. In short, the older a recording is, the lower the probability is for it being played back.

Therefore, a MAID like system is a logical choice in this use case. Another key change from traditional HDD-use is that the Coldstore writes data to disks *sequentially*, meaning that there is no directory area at all. The data is simply appended sequentially, from start to finish, minimizing drive head movement as well as increasing redundancy in case of bit failure.

Veracity exploits the fact that video data itself is recorded sequentially, and thus records it sequentially onto the drives. To increase redundancy of the system, the Coldstore records video onto two disks at the same time. If one of them fails, there is still another disk for backup. When the disks are full, the Coldstore continues to record to the second disk of the pair, and a new disk. This means that each disk will be filled up twice. However, it also means that there is no true redundancy, since the old copy of the data on the disk is overwritten when new data is being written to that disk. The Coldstore also offer a RAID 10-like mode called Full-mirror which, as RAID 10, halves the amount of available data, but offers redundancy in the form of permanent copies of every disk.

When writing to a pair of disks in the standard mode, the Coldstore can deactivate the other drives when they are not being used for recording or playback, meaning that the entire NAS consumes very little power when used, about 60W according to Veracity. Since the drives are powered down most of the time, Veracity also claims that it is possible to load the Coldstore NAS with cheaper lower quality drives, thus bringing down costs even more for the end user. Coldstorlae also features hot swapping, meaning that disks that are not being used for the moment can be replaced while the unit is operating at normal capacity.

The Coldstore is also scalable in two different ways. Either the drives in the device could be emptied when full and replaced with new ones, or more Coldstore devices could simply be added to the system. In the use case relevant for this thesis, either the Colossus or several smaller 3U Pro's could be used to reach the storage requirement.

5.1.3 45Drives

45Drives produces storage servers or NAS:s for large scale applications. The 45Drives Storinator NAS comes in several variants for different sizes as well as different tiers of performance. It also allows for hot swapping of hard drives, meaning that failed drives can easily be removed and replaced without having to worry about disassembly of the system. The system can also continue to run while a hot swap is performed. Since the 45Drives devices are regular NAS:s, they do not offer anything unique in themselves with the exception of the relative cheap price, hot swappable hard drives, and large storage space.

The Storinator comes in different drive sizes, where the smallest one can hold up to 15 drives, and the largest can hold 60 drives. The model that holds 45 drives offer 630TB when using 14TB drives, and is therefore an interesting candidate. Since these devices are regular NAS:s, they need to be configured in RAID mode and, which will be discussed later, this leads to some complications with regards to reduction in effective storage space. For a more in depth discussion on these problems, see Section 6.6.3

5.1.4 Magnetic tape

As mentioned in the storage technologies chapter, magnetic tape has very low price per terabyte compared to other solutions. This made it a natural choice for further examination. For tape to work, a tape drive needs to be directly connected to the server. LTO-8 is the choice of LTO format since it offers the biggest storage spaces and highest throughput. The tapes have to be switched out regularly when they get full, which will include manual intervention. The chosen vendor for this evaluation was HPE.

5.2 Technical evaluation

The products could either be used as primary storage, or as an archival tier storage coupled with a primary storage – located on the server – for more recent recordings. If used as an archival storage solution, performance for the products is not as essential as in the main storage device. The lowest requirement is that the device must be able to handle storing a bitstream of 120Mbit/s. The retrieval and reading time is not essential, hours is considered acceptable.

5.2.1 AWS Glacier

AWS Glacier was tested as an archival solution for older video recordings, since the retrieval times can be very long.

Performance testing

To get AWS Glacier operational requires usage of the provided API (or console) from Amazon. For testing AWS Glacier, a video file was uploaded to the EU North-1 (Stockholm) endpoint. This video file had a size of 102MB, which equals one chunk of video outputted from a single camera. In a real world scenario, there would be several chunks per camera per day.

In AWS, the user can create vaults where files can be uploaded. The vaults can be seen as directories. The files inside vault are called archives, and consists of single files, meaning that it is not possible to upload a folder containing multiple files. If multiple files are to be uploaded as one single upload, they have to be zipped. The upload can be done programatically by using the provided AWS Glacier API, which is available in both C# and Java. The vault is updated once a day, meaning that the changes to the vault are not instantly visible. For EU North-1 (Stockholm), this appears to be around 9PM every night.

To be able to access data from Glacier there are two steps. A retrieval job has to be started. This job can take between 3-5 hours, indicating that there is some form of lower tier storage involved. After this job has been completed, the archive can be downloaded by issuing a HTTP GET request.

To test AWS Glacier's potential, a small program was created. The program uploads the video file from the VMS, to the S3 Glacier EU North-1 (Stockholm) endpoint. It does this by using the .NET SDK for Glacier. The video file was accessed 3 days later from the EU North-1 endpoint. A retrieval job was issued to the Glacier service, which took 3 hours and 43 minutes to complete. After the retrieval job is completed, the file is available to download for 24 hours. This download is very fast since it has been temporarily moved to a hotter tier.

5.2.2 Microsoft Azure

Similarly to AWS Glacier, Microsoft Azure was tested as an archival solution for older video recordings.

Performance testing

To upload files to Microsoft Azure, one can chose to use the command line, SDK, or a web interface. The files that are uploaded are called blobs, and can be stored in different tiers. In contrast to AWS Glacier, the data can be moved to different tiers without changing to another service. When moving data to a different tier, it is inaccessible.

To test the retrieval time of Azure, the same file was used (102MB) as in the AWS Glacier test case. The file was uploaded via Azures web client. The uploaded file (or blob) is initially in the cold tier. The tier was immediately changed to archive tier, which took about 4 minutes. After a week, the file was accessed again. To be able to download the file, the tier has to be changed from archive, since archive files are not readable. This move from archive to hotter tiers is called rehydration, and took about 15 hours. After the rehydration is complete, the file can be downloaded through either of the different Azure tools [14].

5.2.3 Cloud solutions integration into the VMS

After getting an insight into the procedure of storing files in cloud archives it is clear that such a system cannot be used as any kind of primary storage due to the long retrieval times; the VMS would simply be unusable in that case. Instead, the cloud services could serve as a secondary long term storage alternative, where files are uploaded on a regular basis. The more immediate recordings could be saved on a traditional storage device such as the server's hard drive or a connected NAS. Data that is younger than one day is not accessible from AWS Glacier, since it updates the inventory once a day. For Microsoft Azure, it takes time to transfer data into archival tier. If an operator would want to view a very recent recording (from the current day), it would not be accessible if AWS Glacier or Microsoft Azure is the primary storage.

There is however a problem related to picking which data to retrieve from the cloud archival storage. It would be essential to get some glimpse of what the video looks like in a specific time interval, something that is not possible if cloud storage is used. A method that could be used, is to save I-frames (snapshots) at certain times. If there is no motion in the video, the snapshots could have a very long time intervals. If there is much motion, the interval could be shortened. This would give the operator some idea of how the scene looks like around the specified time, before retrieving it from the cloud.

Another method that could be used is to save a low quality stream of the video which takes up very little space. The operator could then judge from this low quality stream if the higher quality stream is needed, and could thus retrieve it from the cloud archive. This is called a *hybrid solution*, since it is a combination of two storage systems, in this case a local cache and a cloud archive.

There would also have to be some kind of native support to download the video files from within the VMS, and it would be preferable if there was a list of ongoing requests and downloads from the cloud services so that the operator can keep track of them. This implementation could require substantial effort to achieve.

Uploading large amounts of data continuously to the cloud

When using cloud solutions, there is the issue with internet speeds. Since the cameras generate 120Mbit/s the upload connection to the cloud provider needs to be faster than this, or otherwise the data will continue to grow before getting uploaded. For example, if the user's upload speed is 20Mbit/s, it would take 6 days to upload a single days worth of recordings. Depending on the camera setup and the user's internet speed, this might not be a problem.

Amazon and Microsoft offer solutions to this problem. They both offer "physical uploading" of data, meaning that they send the user a device into which they can put hard drives. This device is then shipped to the cloud service location and uploaded by the personnel at the site. However, this requires that the user buys several disks, which defeats the purpose of having a cloud solution, and would also introduce a very prominent logistical problem.

5.2.4 Veracity Coldstore

As mentioned before, the Coldstore turns off inactive disks to save power as well as increase the lifetime of the disks. Because of this, it is slower than a standard NAS since the disk where reading currently is done from might be switched off. The spin up time can take 10s of seconds, which leads to less than instant access.

Since the Coldstore operates in a unique way compared to other NAS, the device has to be fully integrated into the VMS if a seamless experience is wanted. The VMS can then let users know that the data they want to playback is on a drive that is currently powered down, and that the retrieval will take some time.

Performance testing

When testing the Coldstore, the device was set as the primary storage for the VMS. The average bit rate for writing was 125Mbit/s to be as representative a testing environment as possible for the Canadian cannabis farms. Higher bitrates were also tested (in the 400Mbit/s range), and the Coldstore could handle them well, indicating that writing performance will not be a problem for our use case.

15 disks were inserted into the device. The disks had varying capacity of 2 and 4TB. The reason for this size was to be able to fill the disks faster and observe the spinning down of disks more frequently.

The retrieval time when accessing data that is stored on an active disk is instantaneous. When accessing data on disks that are turned off, the retrieval time is relatively slow, and there is no indication of what is happening in neither the VMS or the operating system. The retrieval time is in the area of 15-20 seconds or longer on inactive disks. Disks that have been spun up, but are not being recorded to, are kept running for 30 minutes. Since new data is frequently accessed, these 15-20 seconds can be considered to be too long with regards to newer recordings, while being acceptable for older recordings.

Integration with VMS

There are some problems when using the Coldstore without any integration with the VMS. If the Coldstore has recently switched to a new disk for recording, it would mean that even recent recordings would be present on a disk that is powered down. Therefore, as was the case with cloud solutions, a hybrid solution had to be considered. New files are written to the local storage on the server, while older files can be kept on the Coldstore for archival-tier storage. This ensures that recent recordings can be accessed instantaneously.

A program was constructed that moved older recordings to the Coldstore, while keeping the newer recordings on the local storage. The retention time that triggered the move could be changed, but was kept to once every hour for testing purposes. This solution worked well, and new files are instantly accessed from local storage, while older files takes more time (around 15-20 seconds). Recordings in both tiers of storage, hot and cold (see Section 4.3) are still viewable inside the VMS, with the only difference being longer access times when viewing older recordings, i.e. cold recordings. The code from the program was integrated into the VMS and performed well under testing.

5.2.5 45Drives

Since the 45Drives line of products are very traditional devices, and due to time constraints, a storage unit from 45Drives was not tested. The assumption was made that it would perform similar to other NAS:s, therefore a test would not give any new interesting information.

Performance

The 45Drives Storinator line of products offer very good performance since they are standard NAS-devices where all the drives are powered on. The retrieval time for data should be near instant, regardless of the age of the video data. Depending on the RAID level writing speeds can be very high. However, due to how RAID operates, the only RAID levels that are deemed useful are 5, 6 and 10. 45Drives do not recommend more than 15 drives in a RAID 5 or 6 array, since rebuild times can be very long. During rebuild, the throughput of the system is severely degraded and, as mentioned, the drives are more susceptible to failure. Because of this, the only solution is to use several smaller systems of the 15-drive variant, to be able to house the same amount of data, if a system using RAID levels 5 or 6 is desired.

When using RAID 6, the reading is approximately 420MB/s (slowest) and 300MB/s (slowest when writing), and when using RAID 10, the reading is approximately 420MB/s (slowest) and 580MB/s (slowest) when writing [1]. Both of these RAID levels have speeds which are more than adequate for our use case.

Integration with VMS

The different 45Drives models are regular NAS:s, therefore they do not require any kind of additional implementation or integration. There is already support for NAS:s in the VMS as a storage solution. This solution is arguably the easiest out of all the examined products to get operational.

5.2.6 Magnetic tape

Since magnetic tape is such a different technology with regards to what is normally used by Axis, the knowledge and equipment to easily construct and test such a system is lacking. Thus, this implementation will be theoretical in its core, since a real implementation lies outside of the scope of what is possible for this thesis.

A model will be constructed, using the information that has been presented in this report with regards to magnetic tape and its different metrics, and the model will then be evaluated in the same manner as the other systems. The reason for this is to be able to present a fair comparison between all of the different storage technologies.

Performance

The tape drive used in this model will be the HPE StoreEver LTO-8 Ultrium, as presented above, and the tapes will be LTO-8 cartridges from HPE. The model will

be constructed as follows: a site that we are targeting will use up to 60 cameras where each camera has a average bit rate of 2Mbit/s. This amounts up to 1.296TB of video data per day. The LTO-8 cartridge has a capacity of 12TB. This means that one cartridge can store up to 9.2 days of video recordings before it will have to be switched out. The LTO-8 has writing speeds of 360MB/s, so it will take approximately 33 seconds to fill a tape with 12TB of recordings.

Integration with the VMS

The tape drive will not be the primary storage unit, i.e. it will not receive the streamed data directly from the cameras. This means that there is need of a primary storage unit to house the streamed data, before it is backed up in a tape cartridge. This primary storage unit could be the VMS server.

The downside is that the tape drive would have to be directly connected to this external server, which might introduce problems with physical accommodation. The tape drive would not fit inside of the server, and would have to be placed in close proximity.

When the backup to the cartridge is complete, it then needs to be physically removed from the drive and stored in some other manner. This can introduce some problems in how the tape is stored in the physical sense. The operator would have to mark each tape with the timespan of the video with exact dates. The simplest storing method would be to backup each week worth of recording on tape, instead of 9 days. This would nonetheless increase the price per terabyte, since only 7 days would be stored on the tape. 7 days of video data sums up to 9,072TB, which means that almost 3TB would be wasted per tape, which is almost 25%. The payoff would be the increased ease of indexing and storage of the cartridges themselves, since they could be saved on a week by week basis, instead of arbitrary dates. Backup of 9 days would also require the operator to sometimes work on weekends, which might introduce problems. There is also the problem which could occur if the operator misses to backup the data to the tapes, in this case all the data would be lost and overwritten with new data.

There is also the problem with where the tape drive is actually located. As mentioned, it would have to be connected to the server directly. If the client side of the VMS is on a different computer (which would likely be the case), then the operator would have to physically reallocate his person to the area where the server is located and insert or remove the tape cartridge. This might be a problem depending on the setup at each customers site.

And so, to summarize the model: we have a server running the VMS software, a server for storing the streamed data, a tape drive to backup the streamed data once a week to a tape cartridge and a physical index system to store the backed-up tape cartridges. Now some assumptions has to be made with regards to measuring the specific metrics. We assume that, as stated, backup is performed once a week. We assume that the backup-index system allows for retrieval of a specific cartridge in a maximum of 5 minutes. We assume that the distance between the physical index-system and the system required to playback the tape (the servers) is at a maximum distance of a further 5 minutes. Now, to playback the data, it takes, as stated before (5.2.6), "seconds", and so we assume a time of 10 seconds. Since it

Table 5.1: The measured retrieval times for the different systems

System	Retrieval time
Microsoft Azure	15 hours
AWS Glacier	3 hours 46 minutes
HPE LTO-8	6 minutes
Veracity Coldstore	15-20 seconds
45Drives	Near instant

takes time to insert the tape into the drive, we assume the whole operation takes about 1 minute. This gives us an average retrieval time of 6 minutes.

5.3 Summary

In this chapter the different products were evaluated with regards to their technical aspects. See Table 5.1 for a summary of the retrieval times of data.

As stated, 45Drives was not tested in practice since it was assumed that it would perform similarly as other NAS:s'. Both AWS Glacier and Azure Archive were tested by uploading data and retrieving it some time later. It was concluded that the cloud services are too slow to use as primary storage but might be used for archiving older recordings. They would thus require a primary storage unit, such as the server itself to house the data, and then archive older data continuously. The user's internet connection might pose a problem with using cloud, since the upload speed needs to handle large bit rate, which depends on the number of cameras.

The Coldstore can be used as both primary and archival storage, depending on the requirements on retrieval times for recent recordings. The Coldstore was tested as both a primary and archival storage unit, and performed well. The choice was to use a hybrid solution, where older data is moved to the Coldstore for archival storage while newer data is stored on local storage, providing a better user experience.

Magnetic tape was investigated from a hypothetical standpoint, and different implementations were considered. It presents many different shortcomings compared to other storage mediums, and many of the requirements from Axis will be hard to fulfill using magnetic tape.

Economical evaluations

This chapter aims to make an economical comparison and evaluation of the different commercial systems.

6.1 Method

Different technologies and products were evaluated. The metrics used for evaluation were, among others, acquisition costs as well as energy expenditure. The currency used for calculations was USD.

6.1.1 Summary of metrics

The metrics of the economical evaluation of the systems are not only based on acquisition costs – there are other relevant factors to take into account as well. For example, power consumption can lead to expensive energy bills that sum up over time. Failure of storage devices is another cost, and it often manifests itself in the form of a failed hard drive. These have to be replaced by a system integrator, which is also a large expenditure. There is also the case of scalability, meaning how expensive it is to increase the amount of storage.

The basis for the evaluation can be summarized as following:

- Price per terabyte (USD/TB)
- Power consumption
- Operation and maintenance
- Redundancy and failure
- Scalability

6.2 Mathematical model

To compare different solutions, a mathematical model was created, which takes all possible aspects into consideration, and evaluates them quantitatively. The model was made to fit over a five-year timeline, so as to be able to compare the different pricing methods in a fair way. Several assumptions had to be made; for example,

the price per kWh had to be fixated according to Canadian measures [3]. Another assumption that was made was the time to site for the maintenance staff. This was based on the fact that many of the cannabis farms are not located in close proximity to city centers.

Regarding hard drives, after interviewing an Axis employee who worked with storage solutions around seven years ago, the average disk failure rate when using RAID 6 in a system with 60 disks distributed over several storage units, was found to be once per month, i.e. one disk failed every month. Today disks are more reliable and it is thus assumed that the failure rate for disks utilized in a system with 45 disks is about once per four months.

The maintenance time was also set to thirty minutes maximum, as that seems to be standard practice when dealing with these kinds of systems according to Axis themselves, and the fee for the system integrator (the maintenance staff) is set at \$75 per hour. The systems use a maximum of 60 cameras, and the price for each camera is \$418. The server running the VMS software is priced at \$9600. The fee for maintenance is averaged out as a monthly cost for the 5 years of operation. Finally, the amount of data to be stored is averaged out to a maximum of 500TB.

The assumptions can be seen summarized below:

- System uptime is 5 years = 60 months.
- Price for Canadian kWh = \$0.12.
- Time to site for system integrator is set to 1 hour.
- Time to perform maintenance such as switching a faulty drive is set to 30 minutes.
- System integrator fee is \$75 per hour.
- 60 cameras, where each camera costs \$418.
- The price of a server is set to \$9600.
- Amount of data to be stored is 500TB.

For example, if a hard drive in the storage system fails, the system integrator would have to drive to the site and replace it. The time it would take would be 2.5 hours (driving to the site, switching out the drive, driving back). This would amount to about \$188 USD in labour costs.

Note that when using cloud storage services, some of the costs (like the maintenance) amount to 0, and instead costs like the monthly fee for the amount of data used and the cost for retrieval of data are present. Please see earlier sections for a discussion regarding this phenomena.

The model can thus be summarized as follows:

$$TCO = X + Y * (Z_{powerusage} + (Z_{maintenance} * Z_{failurefrequency})) + Z_{procurementcost}$$

and in the case of cloud storage:

$$TCO = X + Y * (Z_{monthlyfee} + (Z_{retrievalfee} * Z_{retrievalfrequency}))$$

where TCO is the *Total Cost of Ownership* in USD, X is the fixed cost of the system (the cameras and the server), in this case $X = 60 * 418 + 9600$, Y is the number of months, in this case $Y = 60$, Z is the specific storage system and Z_m is the specific metric m for the respective costs of the different systems. Note the differences between the cloud storage model and the model for the other systems.

These metrics have been discussed before, and the power usage and the monthly fee are fairly straightforward to calculate. As previously stated, regarding the maintenance, which is connected to the failure rate, some assumptions have to be made. The retrieval fees of the cloud storage systems is a bit harder to estimate; how often would data be retrieved on average during a month, and how much data would be retrieved? To give a fair estimation, we assume that data is retrieved on average four times a year, where each retrieval constitutes one days worth of recordings, i.e. 1.3TB of data.

In the following sections the different metrics will be applied to the different products. In Section 6.9, the TCO will be calculated using the model above.

6.3 Price per terabyte

In this section, the price per terabyte will be calculated for each of the systems, and then compared. For the HDD based systems, the Seagate Skyhawk 14TB was the disk of choice, since it is aimed at video surveillance and is often the most recommended disk of choice for surveillance systems. The Seagate Skyhawk 14TB retails at about \$499.99. The result can be seen in Figure 6.1.

6.3.1 Cloud solutions

The price per terabyte is relatively low for the cloud storage access, at a first glance. The trade off with the low price is that retrieving data comes with a cost that has to be taken into account, and furthermore, that there is a monthly cost for the storage as well, which stacks up over time.

When dealing with the amounts of data required for video surveillance, the price reaches high levels.

AWS Glacier

For AWS Glacier, the price varies from region to region, with US-East being the cheapest at \$0.004 per GB/month, or \$4 per TB/month. For a five year period, this results in \$240 per TB. For 500TB of storage, this price reaches \$2000 per month, or \$120000 over a 5 year period.

In the case of AWS Glacier, they provide different methods of data retrieval. Standard retrievals which usually complete "between 3-5 hours" and which costs "\$0.01 per GB and \$0.05 per 1,000 requests. For example, retrieving 500 archives that are 1 GB each would cost $500\text{GB} * \$0.01 + 500 * \$0.05/1,000 = \$5.025$ ". There are also bulk retrievals, which complete "within 5 - 12 hours" and with a price of "\$0.0025 per GB and \$0.025 per 1,000 requests. For example, retrieving 500 archives that are 1 GB each would cost $500\text{GB} * \$0.0025 + 500 * \$0.025/1,000$ ".

= \$1.2625.” There are also ”expedited retrievals”, which ”typically return data in 1-5 minutes”, and they ”are priced at a flat rate of \$0.03 per GB and \$0.01 per request. For example, retrieving 10 objects with a size of 1GB each, the cost would be $10 * \$0.03 + 10 * \$0.01 = \$0.40$.”[5]

As described in Section 6.2, we assume that cloud data is accessed four times a year, where each access is a download of 1.296TB (one full day). We also assume that one day is stored in one single archive. When using the standard retrieval speed, this would result in the following cost per year, $1296\text{GB} * \$0.01 + 1 * \$0.05/1,000 = \$12.96$. For our case, this would amount to \$51.84 per year.

Microsoft Azure Archival

The pricing for Microsoft Azure Archival varies from region to region, with the US-East price being \$0,002 per GB per month, or \$2 per TB per month. For five years, this price reaches \$120 per TB. For 500TB of storage, the price reaches \$1000 per month and \$60000 for five years, which is half of AWS Glacier.

Microsoft Azure Archive has data writing costs as well, which are \$0.10, per 10 000 API write calls. Also, to be able to download data from the Archive tier (for a discussion about tiers, see Section 4.3), the data needs to be "rehydrated", which means moving it to a hotter tier. This operation costs \$5 per 10000 requests. Downloads also costs \$0.02 per GB [14]. For our use case where data 1.3TB is downloaded four times a year, this amounts up to \$104 per year. Compare this cost to the cost of retrieval in AWS Glacier, which was \$12.96, and it can be seen to be considerably more expensive. In short, storage in Azure is much cheaper than AWS Glacier, but retrieval is much more expensive.

6.3.2 Veracity Coldstore

A Coldstore 3U Pro device has a MSRP of \$7,495.00, while the Coldstore Colossus has a MSRP of \$18,795.00. Using the Seagate Skyhawk 14TB drives that sells for \$499, we arrive at a total price of $499 * 15 + 7495 = \$14970$ for the 3U Pro model, and on the Colossus, $499 * 45 + 18,795.00 = \$41250$. Per TB, this price amounts to \$76/TB for the 3U Pro, and \$67/TB for the Colossus. It should be noted that since the drives are not powered on for most of the time, cheaper and lower quality drives could be used instead of surveillance drives, which would bring down costs even more.

6.3.3 45Drives

The most interesting aspect of the Storinator line of products is the very low cost per terabyte – it is very cheap compared to other alternatives. The price for a Storinator S45, which can hold up to 630TB with 14TB hard drives, is \$9072. A fully loaded Storinator with Seagate Skyhawk 14TB drives costs \$31527. Per terabyte this amounts to about \$50. For the Storinator AV15, the price is \$12856 for a fully loaded model with 15 drives, which is \$61 per terabyte. However, due to RAID performance issues when using a large array, the price per terabyte will not be this low, since not all of the storage can be utilized. This will be discussed in Section 6.6.3.

6.3.4 Magnetic tape

As stated before, the main advantage of magnetic tape is its low cost compared to other storage solutions. A 12TB LTO-8 tape can be bought for as little as \$180. One single LTO-8 tape drive costs around \$5000 [12]. Therefore, tape scales considerably when a lot of data needs to be stored, and is not suitable for smaller amounts.

For using tape, a HPE StoreEver LTO-8 Ultrium 30750 costs \$4,509.99 [12]. If we want 500 TB, we need 42 tape cartridges, where each cartridge costs \$191.99 [13]. For the total cost, we get $191.99 \cdot 42 + 4,509.99 = \12573 . Per terabyte, this comes down to around \$25.

With regards to price per terabyte, there is the concern of the longevity of the format itself. LTO in the past used to be backwards compatible with two generations. It seems however that LTO-8 is only backwards compatible to LTO-7. It is reasonable to assume that this will also be the case with LTO-9 when it is released in the near future. Users will then have to change their tape drives if they want to use the new higher capacity tapes. The new standard will likely reduce the price per terabyte even more, which will increase the benefit of using tape.

6.4 Power consumption

Power consumption refers to how much the storage devices consume in typical usage. There is a difference between typical usage and what the PSU is rated for. To be able to do an unbiased evaluation, we evaluated the products based on their typical power consumption. When calculating costs for power consumption, the price of the Canadian kWh was used, which, as previously stated, at the moment (June 6, 2019) is \$0.12.

6.4.1 Cloud solutions

Since the responsibility of power consumption relies on the cloud provider (as mentioned in Section 3.2, this is referred to as indirect energy), the power consumption is nothing that we can control or calculate, as the technology used is not public. However, since the data retrieval is slow, and the data storage price is lower than that of the regular cloud storage, it is safe to assume that the technology used is cheaper and has lower power consumption than that of the standard cloud storage, i.e. non-archival cloud storage.

6.4.2 Veracity Coldstore

Since only two drives are active at a time (except when performing mass search), the power consumption is very low, only 60W for the 3U Pro model, which houses 15 drives with a total of 196TB. This brings down operational costs as well as environmental impact. The Colossus model with capacity of 630TB only consumes 80W in typical usage, giving it a superior power consumption per terabyte compared to using multiple smaller Coldstore units like the 3U Pro.




Figure 6.1: Price per terabyte, based on product specifications

6.4.3 45Drives

As mentioned, the power consumption for the 45Drives system is relatively high, since all the drives are constantly running. The power supply is rated for 1200W, though under typical use it will be lower. 45Drives themselves have performed tests on the model with 45 drives, and report a typical power consumption of 511W.

The smaller model, the AV15 does not have a reported typical power consumption rate. It is possible to estimate one from the larger model. The 45 sized model uses around 120W for the mainboard CPU, while the rest of the system uses 391W for the drives and other functions, or about 8.69W per drive. The AV15 uses a weaker CPU, and should therefore have a slightly smaller power consumption, 100W should be a reasonable estimate. If every drive uses 8.69W, we arrive at 250W per AV15 device. This gives us an upper bound in our estimates.

6.4.4 Magnetic tape

A major advantage with magnetic tape is that the power consumption is much lower compared to systems using HDDs or SSDs. A tape drive only consumes power when it is being used to write or to read. Since there will only be one tape drive, this amounts to a very low power consumption compared to most solutions that have all the disks running at all times. This gives us the lowest power consumption overall. HPE tape drives use about 24W in typical use [12].




Figure 6.2: The power consumption for the evaluated systems. Note the lack of cloud solutions, since the power consumption of cloud-based systems are unknown to the user.

6.5 Operation and maintenance

The optimal system would let the user forget about the existence of a storage system and just use the VMS. This is not always the case, as some solutions require manual intervention for maintenance.

6.5.1 Cloud solutions

One of the biggest benefits with cloud storage is that there is no maintenance required, since it falls on the cloud provider to carry out any and all maintenance. The user does not have to worry about any operational costs at all other than the actual subscription of the cloud storage itself, which is not classed as a maintenance cost in this report. For operational use, cloud can be fairly easy to use as long as it is implemented well into the system.

6.5.2 Veracity Coldstore

The Coldstore features hot swappable hard drives that can be removed at any time, if for example a hard drive needs to be replaced in the case of failure. New drives are automatically incorporated into the array. It should therefore be very easy to operate. When estimating the required maintenance needed, we had to take into account that Veracity themselves said, as previously stated, that drives in their units would last up to 10 times longer. To be able to perform an analysis as unbiased as possible, we instead went with a lifetime of 5 times longer when calculating maintenance costs. With this in mind, the average time between the need for maintenance is 20 months.

6.5.3 45Drives

The maintenance itself for 45Drives should be straightforward, since it too, as mentioned, features hot swappable hard drives. When hard drives fail they can easily be replaced. It does not require any other form of maintenance in the normal use case where there is no failure of the actual NAS. The frequency of the maintenance needed is another thing, however. As mentioned in Section 6.2, problems arose quite often when using the setup described in that solution — one disk failed every month. Our system will be similar to the system in that description. The failure rate is estimated to be about once every four months when using 45 disks.

6.5.4 Magnetic tape

The major drawback of tape is the intensive manual labour and maintenance involved. The only reasonable solution for a medium sized business is to buy a single tape drive, since auto-loading tape libraries are expensive. This means that only one tape can be used to either write to or read from, at any given time. The result of this is that the tape has to be switched out when it is full of backed up video data, and archived somewhere physically. This physical archiving will also include manual labour where the personnel on site will have to come up with an

archiving solution. A lot of work falls on the user of the surveillance system, which can be a problem.

A positive side effect of this is that data in the form of video surveillance can easily be extracted, since it is located on a single tape or, at most, a few tapes. It can be removed without interfering with the overall system. Tapes are also relatively small and can be transported off-site for backup. For a more in-depth description of how the system would work, see Section 5.2.6.

Another potential drawback from using tape is the maintenance of the drive. All LTO drives requires the use of a cleaning cartridge once in a while to be run in the tape drive. Dirt can build up inside the tape drive and create unrecoverable errors, no matter the cleanliness of the environment it is used in. These cleaning cartridges themselves also have a limited number of uses, further driving up both cost and maintenance [9].

6.6 Redundancy and failure

Redundancy and failure refers to how fault tolerant the storage device is, as well as if it features any kind of fail-safe against data loss. Since missing data might lead to legal repercussions (see Section 1.3.1) such as a loss of the required licences and ultimately bankruptcy for the user, it is important that the surveillance data is not lost and is always available. The problem is that redundancy is expensive, and can in some cases introduce several performance problems as well as big cost increases. The different products provide different solutions for handling redundancy.

Regarding the requirements for redundancy, since all the solutions have different techniques when it comes to tackling the problem, it is difficult to state a general requirement on this aspect. As stated in Section 1.4.3, redundancy is a desirable trait, and the legal requirements imply that 100% of the data from the required period should be available. However, in reality it is impossible to guarantee that no data will ever be lost.

A model revolving around the probabilities of data loss for the different solutions could be constructed. It is however deemed to be outside of the scope of this thesis, and thus, the requirement in this case is that the redundancy is simply desired to be as good as possible.

6.6.1 Cloud solutions

Since the responsibility of redundancy falls on the cloud provider when using a cloud service, the end user does not have to worry or take this into account due to the cloud storage service providers, in this case both Amazon and Microsoft, often utilizing different forms of redundancy, from local redundancy to geo-redundancy [7][6]. The cost of redundancy and failure is therefore already included in the price of cloud storage.

6.6.2 Veracity Coldstore

As mentioned, the drives that are used in the Coldstore products are powered off most of the time, and so the risk of failure is lower than in a traditional "always-on"

system. Since the data is being written sequentially on the drive, it means that there is no RAID mode involved in the process. On a hard drive in typical usage, if the directory area is corrupted, the whole array is lost. On a hard drive used in a Coldstore unit, there is no directory area, so there is no risk of large data loss. Also, different RAID setups allow for different degrees of failure. For example, if three drives fail on a RAID 6 array, the whole array is lost. On the Coldstore, only the data on the failed drives is lost.

This also means that it is possible to remove several drives from the system, which if done on a RAID 6 array would crash the entire system.

If a drive failure occurs during writing, there is another drive where a copy of the data is written to. The faulty drive can be removed and replaced. If the other drive would also fail at this point in time, there is no error recovery.

However, there is still the risk of drive failure after the drive has been turned off. If the operator wants to view the recordings on the disk, the drive is spun up again. If a drive failure was to occur at this moment, all the data on that disk would be lost, and there is no other redundancy in place to backup the data. Veracity themselves claims that they have never heard of this occurring in the Coldstore's deployment lifespan, even when a unit was using disks that Veracity themselves refer to as being of lower quality. Nonetheless, it is still important to note that there is no redundancy that can restore the content on a failed drive.

As mentioned, the Coldstore also features a Full-mirror mode, meaning that all of the data will be stored on two disks, for true redundancy. This closely resembles RAID 10. The downside of this mode is of course that the capacity is halved, from 196TB to only 98TB on the 3U Pro unit, for example.

6.6.3 45Drives

45Drives offers RAID for redundancy. For RAID, the levels available are 0, 5, 6, 10. 45Drives themselves do not recommend using anything other than RAID 10 or 0 when 15 or more drives are being used because of long rebuild times. After talking with Axis employees, we came to the same conclusions. This drastically changes the price per terabyte for this unit.

Using RAID 10 in the device halves the effective storage space, which is a big concern as it doubles the price per terabyte. Using RAID 6, smaller Storinator units have to be used - the Storinator AV15 houses 15 disks and can thus be used in RAID 6 mode, but then a larger amount of storage devices has to be used to reach the required amount of storage needed, further driving up the price per terabyte. The power consumption will also increase since more units are required.

6.6.4 Magnetic tape

Because of the nature of magnetic tape, there is no built in redundancy in the format. If only the minimum number of required tapes are used, it means that there is no backup of the data at all. To introduce redundancy, the data will have to be mirrored onto another tape cartridge, which will bring up the cost as well as the labour involved. Instead of switching out the drive once a week, the operator will need to insert a cartridge, copy the data, then insert the second cartridge

and copy the data again. This will also increase the price per terabyte since the number of tapes involved will double. It will also introduce larger physical storage requirements as in where to house the approximately hundred tapes, as well as complexity with regards to how to properly index them, and finally, cleaning will also have to be performed on twice the amount of cartridges.

6.7 Scalability

Scalability is an important metric to take into account when choosing a new storage system. Scalability refers to the capacity for a system to change in size. If new laws are implemented that extend the data retention, or in some other way increase the amount of storage needed for the solution, it needs to be able to handle the new requirements. There are different forms of scalability, as discussed in Section 3.4, and they have different costs related to them.

6.7.1 Cloud solutions

Scalability is one of the main benefits of cloud storage, since no hardware is required. More storage can simply be bought and used by the user.

6.7.2 Veracity Coldstore

To handle increasing storage demands, many Coldstore units could be used. The Coldstore also has the benefit of allowing the user to remove the drives and insert new ones. Therefore one Coldstore unit could still be used if data retention laws introduce new requirements. If the user still wants all of the data readily available (without having to manually insert drives into a Coldstore unit), more Coldstore units can be bought, i.e. the Coldstore utilizes so-called horizontal scaling.

6.7.3 45Drives

In contrast to the Coldstore, since 45Drives operates in RAID mode, the disks can not be removed and inserted later into the device to read from them. This will trigger a rebuild and will take a considerable amount of time. Therefore, to gain more storage space, a new Storinator unit has to be purchased. The Storinator units thus scale in the same way as the Coldstore units, if several Coldstores are used instead of archiving the drives. This form of scaling can be visualized in graph 6.5, where each new unit represents a "bump" in the graph - see for example the yellow line, representing a 45Drives Storinator S45 in RAID 10 mode. As stated, 2 or more units are required to reach the storage needed (500TB) for this use case.

6.7.4 Magnetic tape

One of the major advantages with magnetic tape is that it scales incredibly well. The drive is the most expensive investment, with the tape cartridges being cheap in comparison to SSDs and HDDs. No new drive is needed if retention times were

to increase, more cartridges can simply be bought and used — so called vertical scaling.

6.8 Price per TB and power consumption revisited

With these evaluations in mind, it is worth revisiting the \$/TB and power consumption, since taking redundancy and scalability into consideration will bring the costs up for the physical systems - Veracity Coldstore, 45Drives Storinator and the magnetic tape system - with regards to these aspects.

See graph 6.3 for the revisited graph for \$/TB. Note that more units are needed for the Coldstore- and 45Drives' products, but for the magnetic tape system, we simply need more tapes to back up the data on and do not need to use more tape drives. This is because – as previously mentioned – tape scales differently. Furthermore, which will be discussed in chapter 5, the tape drive will not function as the primary storage unit, but will rather act as a storage for the *archived data*, i.e. the long term stored data.

Also note how there are two more columns showing the Full-mirror mode of the Coldstore units. Full Mirror-mode is, as mentioned, the same as RAID 10, but for the Coldstore line of products, thus utilizing the LAID system described earlier. Two units of each are needed to make a fair comparison to the other cases, since RAID 10 halves the storage capacity.

It is important to note that the Full Mirror mode is not required to achieve redundancy due to the Coldstore not using RAID in the first place (see Section 5.1.2). Even if several disks were to fail in a Coldstore unit it does not amount to the same disastrous consequences as several disks failing in a RAID system. A Coldstore unit can lose several disks, while three disks failing in RAID 6, for example, destroys the whole array.

As can be seen, for some solutions, the cost goes up quite a bit. The most obvious increases come from the 45Drives Storinator units since they are in RAID mode; RAID 6 and RAID 10 for the AV15 and S45 units respectively. As has been mentioned, RAID is required for there to be any redundancy in traditional systems.

With regards to the power consumption, this cost also goes up quite a bit when taking the full scale of the solutions, required for the storage capacity and the redundancy to be acceptable, into account for some of the systems. See graph 6.4. For example, to reach the required amount of storage when using Coldstore 3U Pro, 3 units are needed, which effectively triples the power consumption. The same can be seen with the 45Drives systems, the AV15 and the S45.

The only physical storage units which did not see an increase in power consumption when looking at the complete solution were the Coldstore Colossus, with its 616TB of effective storage, and the magnetic tape drive, since it scales vertically and simply requires more tape cartridges for there to be redundancy, thus requiring no extra units to be running.




Figure 6.3: The revisited price per terabyte, with redundancy and scalability factored in




Figure 6.4: The revisited power consumption, with redundancy and scalability factored in. Again the cloud-based systems are not shown, since the power consumption of those solutions are unknown to the users.

6.9 Evaluation

In this section, all of the products will be evaluated with regards to the mathematical model described in Section 6.2. All the aspects will be taken into account, giving a total cost of ownership for each solution, and the products will be compared with each other.

In table 6.1, the TCO for all of the products can be seen.

6.9.1 45Drives

First off, there are many different configurations available for the 45Drives devices. Since RAID 5 and 6 are not recommended in larger systems, they are excluded for the model using 45 drives. We are left with 4 choices, either the S45 model in RAID 0 or 10, or the AV15 model in RAID 0 or RAID 6. As mentioned, RAID 0 does not offer any redundancy - if one drive fails, the whole array fails. It therefore poses a huge risk, one that is not acceptable, and so RAID 0 has to be discarded. It is however still good to make a comparison between RAID 0 and other RAID levels.

In Figure 6.5, the total cost per terabyte can be seen for the different 45Drive models. As the graph shows, RAID 10 for the S45 model is very expensive since the effective storage is cut in half, but offers very good performance with no rebuild time penalties. Please note that this is not a graphical representation of the TCO, i.e. power expenditures, costs for cameras and maintenance etc., are not taken into consideration, but only the actual costs of the units when a specific amount of data is wanted.

Total Cost of Ownership

For the different model setups, there are different TCOs. Since, as mentioned, using RAID 10 requires double the amount of hard drives, it will of course be a more expensive option. When estimating failure rate, the assumption was made that a drive fails every 4 months in a system with 45 drives, and every 2 months in a system with 90 drives.

When using the mathematical model for TCO as described in Section 6.2, we arrive at what can be summarized in graph 6.6a for the AV15 model in RAID 6. Respectively, the TCO for the S45 model in RAID 10 can be seen in graph 6.6b.

6.9.2 Veracity Coldstore

Due to the fact that the Coldstore units do not use any RAID level, they can utilize more storage from the disks than that of a conventional RAID system, which gives them a lower total cost. For a comparison between the different units, see Figure 6.7. Again, please note that this is not a graphical representation of the TCO, but only the actual costs of the units themselves with when a specific amount of data is wanted.

45Drives RAID level comparison




Figure 6.5: Comparison between different Storinator units and configurations


Total cost of ownership

For the Coldstore line of products there are, as stated, two different units available - the Colossus and the 3U Pro. For our storage requirement of 500TB, one Colossus is sufficient while three 3U Pro units would be needed.

The TCO equation (6.2) was used. The frequency of maintenance was set to once every 20 months. The power consumption for the Coldstore 3U Pro is on average 60W, and 80W for the Colossus unit. These numbers produce the TCO:s that can be viewed in Figure 6.8a and 6.8b for the 3U Pro and the Colossus, respectively.

6.9.3 Cloud storage

Even though cloud storage offers very low prices per terabyte, it is a monthly fee that quickly sums up; the price stacks up over time, and will eventually be more expensive than having a physical storage unit in-house. It is however very accessible and easy to use. There is no maintenance required, and the demand on redundancy falls on the cloud provider, cutting down the cost for physical hardware.



(a) The TCO for Storinator AV15 in RAID 6 mode. 3 devices are needed to be able to store the required amount of data. (b) The TCO for Storinator45 in RAID 10 mode. 2 devices are needed to be able to store the required amount of data.

Figure 6.6: TCO:s for the Storinator units (in USD)

Veracity Coldstore comparison (Full Mirror is equal to RAID 10)



Figure 6.7: Comparison between different Coldstore units and configurations

Total cost of ownership

Nonetheless, the price can be high, especially for AWS Glacier, which is the most expensive solution out of all the solutions that were investigated. The TCO for AWS Glacier can be viewed in Figure 6.9a


The Azure archive solution provides a good total price over five years which can compete with the physical devices. Cloud solutions offer total redundancy as well as geo-redundancy, something that the physical devices do not. This makes Azure Archive very competitive, and the TCO can be viewed in Figure 6.9b.

6.9.4 Magnetic tape

Magnetic tape is a very good alternative with regards to economical aspects. It provides the cheapest price per terabyte compared to all other solutions. Furthermore, the power expenditure is very low. It scales very well, since all that is required to get more storage is to buy more cartridges.

Total cost of ownership

The TCO for a system consisting of a backup tape drive can be seen in Figure 6.10.




(a) The TCO for Coldstore 3U Pro. 3 devices are needed to be able (b) The TCO for Coldstore Colossus. One device is needed to be able to store the required amount of data.


Figure 6.8: TCO:s for the Coldstore units (in USD)

AWS Glacier

TCO overview



■ Devices ■ Storage ■ Electricity ■ Yearly Software & Fees
■ Data retrieval




Total Cost of Ownership 155788



(a) The TCO for AWS Glacier

Azure Archive

TCO overview



■ Devices ■ Storage ■ Electricity ■ Yearly Software & Fees
■ Data retrieval



Total Cost of Ownership 96248

(b) The TCO for Microsoft Azure Archival

Figure 6.9: TCO:s for the cloud service systems (in USD)

HPE LTO-8




Figure 6.10: The TCO for HPE LTO-8, magnetic tape (in USD)

Table 6.1: The TCO:s for all of the evaluated systems in descending order of cost. The amount of data to be stored in this example is 500TB and the time-period is 5 years.

System	TCO (USD)
AWS Glacier	155,788
Coldstore 3U PRO Full Mirror	134,734
Storinator S45 RAID 10	125,465
Coldstore Colossus Full Mirror	123,282
Microsoft Azure Archive	96,248
Storinator AV15 RAID 6	87,284
Coldstore 3U PRO	85,229
Coldstore Colossus	79,503
HPE LTO-8	55,835

6.10 Discussion

The TCO for all of the evaluated systems can be seen in table 6.1.

The systems from 45Drives and Coldstore are the ones that offer the lowest price per terabyte if tape is excluded. The clear advantage of the Coldstore is how it handles the disks, where lifetime can be much longer than in a NAS device utilizing RAID, thus minimizing the costs for maintenance and disk failure. Veracity themselves promise up to 10 times the lifetime compared to when being used in a conventional device. This significantly brings down operational costs.

The power consumption is also much lower, not only because the Colossus solution only requires one single device, but also because the power consumption per unit is much lower. The total storage space of a fully loaded Colossus is also larger than that of three Storinator AV15-units (616TB vs 546TB), since the Colossus does not utilize RAID.

The cloud storage services prevail in the area of redundancy and power consumption, where the responsibility falls on the service provider. As of now, there is no economical benefit of using AWS Glacier over Microsoft Azure Archive – the difference lies in the longer retrieval time.

In short, there is not much difference with regards to price when comparing Coldstore Colossus to 45Drives. AWS Glacier is expensive when compared to the other solutions, while Microsoft Azure Archive provides very good price per terabyte over a five year period.

6.11 Summary

In this chapter, the different commercial systems chosen for evaluation were presented and evaluated based on several economically-related criteria that summarized acquisition costs with operational costs into a total cost of ownership. Economically speaking, tape is very cheap per terabyte. It was concluded that Verac-

ity Coldstore is cheaper than 45Drives when including aspects of failure rate and maintenance. For cloud solutions, the price for Azure Archive is much cheaper compared to AWS Glacier.

Conclusion

With all the evaluations in mind, different conclusions can be reached with regards to what metric the user finds most important. In the end, a conclusion was reached as to what can be considered to be an answer to the question of long term storage and how to best perform it in the context of this thesis.

If the only metric of interest is the final cost, then magnetic tape is the best solution. However, as has been stated many times previously, the tape-solution will most likely cause more problems than the other solutions with regards to the manual labour involved. It is also more complicated to integrate into the current VMS, and more consideration has to be taken to the physical accommodation of the system. It is also likely that tape will not fulfill some of the requirements required for the VMS to function with all of its current features. Thus, a decision has to be made with regards to if the low price justifies these shortcomings.

If the most important metric is redundancy, then Microsoft Azure is the most dependable choice. Azure is chosen over AWS Glacier since it is cheaper, but both of the cloud solutions offer redundancy that simply cannot be matched by any of the other solutions, in the form of data being backed up several times — at multiple places even — thus protecting it from potential danger in the form of disk failure and other hazards such as fire outbreaks.

With regards to operation and maintenance, then the cloud systems were again superior, since the user does not have to perform any maintenance at all, once the VMS is configured to send the data to be archived to the cloud service. Furthermore, with regards to scalability, the cloud solutions also outperform the other options — as stated, the user can simply purchase more storage if and when the need arises.

The best price per TB was also achieved by the cloud solutions, Microsoft Azure specifically, but as has been mentioned before, this metric can be a bit misleading since the user pays a price per month for the cloud systems. If we consider a system that requires 500TB of storage yearly, over a period of 5 years, as has been the use case in earlier sections, then the best price per TB is actually achieved by the tape-solution. It is, however, possible that paying the price for acquiring one of the other systems can be too high for the end user. Cloud systems could then become preferable due to their monthly bill, even if the sum of the monthly payments become higher after a couple of months or years.

The lowest power consumption was achieved by magnetic tape. But, as mentioned, that solution was deemed as unsuitable due to other factors.

The best retrieval times were achieved by the Storinator units. These units function as normal NAS:s, where the disks are always active and the data is thus instantaneously accessible. The system which would be the easiest to implement, that is, integrate into the current VMS, would be a regular NAS-system, and with regards to the investigated systems it would be one of the Storinator units.

The best solution, if the metrics are weighted against each other, with price still being the most important one and with magnetic tape discarded, is according to the authors the Veracity Coldstore; more specifically, the Coldstore Colossus. As stated, with magnetic tape discarded, it is the cheapest solution (see Table 6.1). Furthermore, only one unit is required, as opposed to up to three units in some of the other solutions. The retrieval times are practically instantaneous if the data to be accessed is on the currently recording disk (in this case, it functions like a NAS), and is at most up to 20s if the data is on an idle disk and the disk thus needs to be spun up. If a hybrid solution is used, such as the one constructed in the testing Section (Section 5.2.4), then data that is newer is always instantaneously available.

The power consumption is incredibly low if compared to other systems with similar storage capacity, like the Storinator line of products, due to the spin-down technology mentioned earlier. The maintenance required for a Coldstore Colossus, if the unit is properly integrated (see Section 5.2.4), is also close to zero, and it should only ever require maintenance if a disk were to fail.

And as such, the only aspect in which the Coldstore Colossus is really lacking, is redundancy. There is no real redundancy for the disks other than for the one that is currently being recorded to – as explained earlier, it is mirrored by a separate disk that is recording at the same time. But the disks that are full have no other backup, if they were to fail, that data is lost. However, it is very unlikely that they would fail, since they are not powered on when they are not in use. Also, a failing disk will never affect any other data than that which is stored on the failed disk itself.

With regards to scalability, the Colossus also performs quite well. The current maximum storage space for a Colossus unit, utilizing 14TB drives, is 616TB, which is a bit more than the 500TB that was used in the use case. Furthermore, 500TB is also a little bit above what is actually required in terms of storage space (see Section 1.4.2), since that number was chosen to be able to have a margin in the calculations. If the required storage space drastically increases, it is possible to acquire disks with a larger individual capacity (which should be available in the near-future, see Section 3.4) to increase the storage capacity even more.

In summary, the Coldstore line of products differed from other NAS:s in how they work (sequential recordings, spinning down of inactive and idle disks – see Section 5.1.2), which was interesting. However, the possibility of the Coldstore performing unexpectedly due to these differences was still present, and so it was important to test it.

As stated in Section 5.2.4, the Coldstore performed well under testing. It was integrated into the VMS without any issues and this, together with the other advantages mentioned above, made it our final recommendation.

References

- [1] 45Drives specification. Accessed on: 2019-02-19. URL: <https://www.45drives.com/wiki/index.php?title=Configurations>.
- [2] A typical RAID 10 configuration. Wikimedia Commons. Used with permission under Creative Commons CC0 1.0 Universal Public Domain Dedication. Accessed on: 2019-05-29. URL: https://commons.wikimedia.org/wiki/File:RAID_10_01.svg.
- [3] Average electricity prices around the world: \$/kwh. Accessed on: 2019-02-20. URL: <https://www.ovoenergy.com/guides/energy-guides/average-electricity-prices-kwh.html>.
- [4] AWS & Sustainability. Accessed on: 2019-03-10. URL: <https://aws.amazon.com/about-aws/sustainability/>.
- [5] AWS Glacier. Accessed on: 2019-01-25. URL: <https://aws.amazon.com/glacier/faqs/#dataretrievals>.
- [6] AWS Global Infrastructure. Accessed on: 2019-02-20. URL: <https://aws.amazon.com/about-aws/global-infrastructure/>.
- [7] Azure Storage redundancy. Accessed on: 2019-03-20. URL: <https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>.
- [8] Cannabis Regulations SOR/2018-144. Accessed on: 2019-05-02. URL: <https://laws-lois.justice.gc.ca/eng/regulations/SOR-2018-144/FullText.html>.
- [9] Cleaning your tape drives. Accessed on: 2019-03-25. URL: https://www.ibm.com/support/knowledgecenter/en/ssw_ibm_i_71/rzam4/rzam4cleantape.htm.
- [10] Diagram of a RAID 6 setup, which is identical to RAID 5 other than the addition of a second parity block. Wikimedia Commons. Used with permission under Creative Commons CC0 1.0 Universal Public Domain Dedication. Accessed on: 2019-05-29. URL: https://commons.wikimedia.org/wiki/File:RAID_6.svg.
- [11] How Can You Prevent Magnetic Tape from Degrading Prematurely? Accessed on: 2019-03-19. URL: https://www.clir.org/pubs/reports/pub54/5prematre_degrade/.

-
- [12] HPE - Small Business Marketplace. Accessed on: 2019-02-25. URL: https://buy.hpe.com/pdp?prodNum=BC022A&country=us&locale=en&catId=12169&catlevelmulti=12169_304612_3446236_4150338.
- [13] HPE - Small Business Marketplace. Accessed on: 2019-02-25. URL: https://buy.hpe.com/pdp?prodNum=Q2078A&country=us&locale=en&catId=12169&catlevelmulti=12169_64242_342326_34648.
- [14] Microsoft Azure Archive. Accessed on: 2019-03-10. URL: <https://azure.microsoft.com/sv-se/pricing/details/storage/blobs/>.
- [15] Open Grid Scheduler / Grid Engine. Image rectification. Wikimedia Commons. Used with permission under Creative Commons CC0 1.0 Universal Public Domain Dedication. Accessed on: 2019-01-20. URL: <https://en.wikipedia.org/wiki/File:HuaweiRH2288HV2.JPG>.
- [16] Seagate Continues to Lead as HAMR Technology Advances. Accessed on: 2019-03-10. URL: <https://blog.seagate.com/business/seagate-continues-to-lead-as-hamr-technology-advances/>.
- [17] Storage Area Network (SAN). Accessed on: 2019-03-20. URL: <https://www.techopedia.com/definition/1116/storage-area-network-san>.
- [18] StorageTek Powderhorn tape libraries. Wikimedia Commons. Used with permission under Creative Commons Attribution-Share Alike 2.0 Generic. Accessed on: 2019-02-10. URL: https://en.wikipedia.org/wiki/File:StorageTek_Powderhorn_tape_library.jpg.
- [19] Supertape data storage capacities. Wikimedia Commons. Used with permission under Creative Commons Attribution-ShareAlike 3.0. Accessed on: 2019-01-25. URL: https://en.wikipedia.org/wiki/File:Supertape_data_storage_capacities.svg.
- [20] The State Of Cannabis Video Surveillance Requirements In CA, OR, WA, AK And Canada. Accessed on: 2019-05-02. URL: <https://www.arxys.com/cannabis-video-surveillance-requirements/?fbclid=IwAR2Vc6HpZTpqR0D-zBX4nevCaI6udCD0aL4vXoccaGOVuc5g1UYYYzYrcsg>.
- [21] Understanding scalability in data storage. Accessed on: 2019-03-02. URL: <https://stonefly.com/blog/understanding-scalability-in-data-storage>.
- [22] Paul Acorn. Seagate Plans To HAMR WD's MAMR; 20TB HDDs With Lasers Inbound. Accessed on: 2019-03-15. URL: <https://www.tomshardware.com/news/seagate-wd-hamr-mamr-20tb,35821.html>.
- [23] Antony Adshead. Microsoft revives MAID with Pelican, but tape can still sleep easy. Accessed on: 2019-02-20. URL: <https://www.computerweekly.com/blog/StorageBuzz/Microsoft-revives-MAID-with-Pelican-but-tape-can-still-sleep-easy>.
- [24] Harrison Astbury. The real cost of running your fridge. Accessed on: 2019-03-25. URL: <https://www.canstarblue.com.au/appliances/how-much-energy-does-a-refrigerator-use/>.

-
- [25] Ali Ismail Awad and Bilal Charif. Business and Government Organizations' Adoption of Cloud Computing. In *Intelligent Data Engineering and Automated Learning – IDEAL 2014: 15th International Conference*, 2014.
- [26] Krish Bandaru and Kestutis Patiejunas. Under the hood: Facebook's cold storage system. *Facebook Code*, may 2015.
- [27] Roderick Bauer. What's the diff: Hot and cold data storage. *Backblaze*, february 2019.
- [28] Brian Beach. Hard Drive Temperature – Does It Matter? *Backblaze*, may 2014.
- [29] P.M. Chen and D.A. Patterson. Storage performance-metrics and benchmarks. *IEEE*, august 1993.
- [30] Dennis Colarelli and Dirk Grunwald. Massive Arrays of Idle Disks For Storage Archives. *Proceedings of the IEEE/ACM SC2002 Conference (SC'02)*, 2002.
- [31] Cosmos Darwin. Understanding SSD endurance: drive writes per day (DWPD), terabytes written (TBW), and the minimum recommended for Storage Spaces Direct. *Microsoft Blog Technet*, august 2017.
- [32] Anna K Eriksson. Så mycket el drar Facebook i Luleå. *Svenska Dagbladet*, jun 2013.
- [33] S. Balakrishnan et. al. Pelican: A building block for exascale cold data storage. *Microsoft research*, 2014.
- [34] Larry Freeman. What is old is new again. PowerPoint presentation by NetApp. Accessed on: 2019-02-14. URL: http://www.snia.org/sites/default/education/tutorials/2012/spring/storman/LarryFreeman_What_Old_Is_New_Again.pdf.
- [35] Hemant Gaidhani. Speeds, Feeds and Needs – Understanding SSD Endurance. *Western Digital Blog*, january 2015.
- [36] James Glanz. Power, Pollution and the Internet. *The New York Times*, sep 2012. <https://www.nytimes.com/2012/09/23/technology/data-centers-waste-vast-amounts-of-energy-belying-industry-image.html>.
- [37] Kazutoshi Katayama, Yuka Chinda, Osamu Shimizu, Tatsuo Mikami, Mayumi Suzuki, and Hitoshi Noguchi. Long-term stability of magnetic tape for data storage under an accelerated condition. *IEEE Transactions on Magnetics*, 52, july 2016.
- [38] Patrick Kinnerk. To Geo or Not to Geo...? Industry Perspectives - Data Center Knowledge. Accessed on: 2019-02-25. URL: <https://www.datacenterknowledge.com/archives/2017/01/26/geo-not-geo>.
- [39] Benjamin Kraft. SSD-Langzeittest beendet: Exitus bei 9,1 Petabyte. *Heise Online*, 2017.

-
- [40] Randy Krum. The lifespan of storage media. Accessed on: 2019-02-08. URL: <https://coolinfographics.com/blog/2012/8/14/the-lifespan-of-storage-media.html>.
- [41] Santosh Kumar and Rajesh Vijayaraghavan. *Solid State Drive (SSD) FAQ*. Dell. Accessed on: 2019-02-25.
- [42] N. Mandagere, J. Diehl, and D. Du. Greenstor: Application-aided energy-efficient storage. In *24th IEEE Conference on Mass Storage Systems and Technologies (MSST 2007)*, pages 16–29, Sep. 2007. doi:10.1109/MSST.2007.4367961.
- [43] Rich Miller. Inside Facebook’s Blu-Ray Cold Storage Data Center. *Data Center Frontier*, jun 2015.
- [44] Bradley Mitchell. An in-depth guide to the differences between san and nas. Accessed on: 2019-01-19. URL: <https://www.lifewire.com/san-vs-nas-818005>.
- [45] Tom Nelson. Do HDDs or SSDs Need ‘Exercise’? The Rocket Yard Investigates. Accessed on: 2019-02-20. URL: <https://blog.macsales.com/43702-we-bet-you-didnt-know-that-your-hdds-or-ssds-may-need-exercise-too>.
- [46] Michael Nuncic. How Long do SSDs Really Last? Accessed on: 2019-02-25. URL: <https://www.ontrack.com/blog/2018/02/07/how-long-do-ssds-really-last/>.
- [47] Overland-Tandberg. *Overland-Tandberg LTO Ultrium Data Cartridges & Cleaning Cartridges*. Accessed on: 2019-03-19. URL: http://www.tandbergdata.com/default/assets/File/Data_Sheets/DS_LTO_media_US.pdf.
- [48] Riley Panko. iCloud Security Survey 2017: Is It Safe? Accessed on: 2019-02-30. URL: <https://clutch.co/cloud/storage/resources/icloud-survey-2017>.
- [49] Riley Panko. iCloud Security Survey 2017: Is It Safe?, march 2017. URL: <https://clutch.co/cloud/storage/resources/icloud-survey-2017>.
- [50] David A Patterson, Garth Gibson, and Randy H Katz. Case for Redundant Arrays of Inexpensive Disks (RAID). *SIGMOD '88 Proceedings of the 1988 ACM SIGMOD international conference on Management of data*, 1988.
- [51] David A. Patterson and John L. Hennessy. *Computer Organization and Design: The Hardware/Software Interface*. Elsevier, 2014.
- [52] David Pease, Arnon Amir, Lucas Villa Real, Brian Biskeborn, Michael Richmond, and Atsushi Abek. The linear tape file system, 2010. URL: <https://ieeexplore.ieee.org/document/5496989>.
- [53] Eduardo Pinheiro, Wolf-Dietrich Weber, and Luiz Andre Barroso. Failure Trends in a Large Disk Drive Population. *Google Inc*, 2007.

-
- [54] A.H.M.E. Reinders, K. Vringer, and K. Blok. The direct and indirect energy requirement of households in the european union. *Energy Policy*, january 2003.
- [55] Andrea C. Arpaci-Dussea Remzi H. Arpaci-Dussea. *Operating Systems: Three Easy Pieces*. University Of Wisconsin–Madison, 2014.
- [56] Matthew Rogers. How Long Can a Hard Drive Hold Data Without Power? Accessed on: 2019-02-20. URL: <https://lifehacker.com/how-long-can-a-hard-drive-hold-data-without-power-5808858>.
- [57] Sriram Sankar, Mark Shaw, Kushagra Vaid, and Sudhanva Gurumurthi. Data-center scale evaluation of the impact of temperature on hard disk drive failures. *Trans. Storage*, 2013.
- [58] Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Hermlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. United States Data Center Energy Usage Report. *Federal Energy Management Program of the U.S. Department of Energy*, june 2016.
- [59] Steve Shen. Vibration of hard disk drives. Accessed on: 2019-02-10. URL: https://www.me.washington.edu/research/faculty/ishen/hdd_vibration.
- [60] Mark W. Storer, Kevin M. Greena, Ethan L. Miller, and Kaladhar Voruganti. Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage. In *Conference on File and Storage Technologies*, 2008.