LONG TERM DATA RETENTION

By

ERIC FIZUR

A thesis submitted to the

Graudate School-Camden

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Masters of Science

Graduate Program in Computer Science

Written under the direction of

Dr. Jean-Camille Birget

And approved by

____

Dr. Jean-Camille Birget

____

Dr. Suneeta Ramaswami

____

Dr. Sunil Shende

Camden, New Jersey

January 2020

THESIS ABSTRACT

Long Term Data Retention

by ERIC FIZUR

Thesis Director:
Jean-Camille Birget

Most information today forgoes a solely physical medium and resides in a digital format; however, the information may be altered or lost over time. There is a need to create a library that will persist for generations with relevant information and be accessible to anyone globally. Possible formats for data storage include microfilm, magnetic disks, solid state drives, DNA data storage, optical data storage, holographic data storage, and cloud computing. All give various solutions to longevity and physical and data integrity but the proposed route utilizes cloud computing due to its growing market and increase of use in businesses and education. The data contained within needs to maintain a high level of data authenticity and integrity. It will require a way to maintain a system of error-correcting codes to make sure the data is unaltered during storage or transfers. I discuss the basic architecture of data centers, the cost of powering them, utilizing more space as data is added, and global load balancing. In regards to the data, I address containment, and solutions on how to deal with environmental and human accidents, severe weather conditions, law and copyright issues, hacking and in extreme cases an electromagnetic pulse from a nuclear explosion. The main data centers should be duplicated in distant and secure locations in both developing and established areas. Data contained in the library must be analyzed for relevance and review to prevent additions of unnecessary or inaccurate data. Finally, I propose best practices with the current information compiled on the creation of a long-term data retention library and what possible future solutions instead.

Contents

Chapter 1

Introduction

A library is a fantastic source of information about the past and relative present. However libraries do not last forever. Chances are your collection of physical monographs that are dwindling if they have not disappeared entirely. Their disappearance is an unfortunate circumstance not for a lack of purpose but due to obsolescence. We need to pass our history to future generations through whatever means necessary. Civilizations transcribe information through various mediums like stones and vases, scrolls and paper and more currently to floppy disks and flash drives. However, like most materials they are lost over time either due to natural or man-made disasters or the material itself losing integrity. Paper books are still prevalent today but for how much longer with data becoming digital? In a similar vein, are cassettes or microfilm still relevant? At the time of this creation cassettes are long since obsolete and microfilm usage has decreased significantly. Though we have learned much from ancient texts and materials, it can still be said that most have been lost over time. Civilizations like the Olmecs are almost completely unknown from the lack of surviving transcriptions of any medium. Imagine what more could have be known if we had concrete information from 500 or 1000 years ago. What if we had accurate measurements of sea levels dating back thousands of years? The information could be useful in finding patterns and creating predictions today. We have a duty to safeguard our known history and data for future generations. And unlike Wikipedia which any and everyone can manipulate, it is necessary to have a database of accurate and authoritative information. The information stored should not last only 5, 10, or 50 years but hundreds or thousands so that accurate information is held for generations.

This paper will discuss the format issues of the current day, the need of authenticity and integrity, and purposeful and accidental damaging situations that may be encountered. Then it will review the requirements on the architecture, costs and locations and methods to add entries into the library and then coalesce all the ideas together to provide possible solutions for long term data retention, with current and future technology in mind.

Chapter 2

Failures

2.1 Obsolescence

It is difficult to predict brand new technology, though one can assume current technologies will improve or produce an alternative way of creating retention. Current technology will in time become obsolete and the data stored unable to be read, the data carrier itself will decay and in time must be changed. This is obsolescence, the process of becoming unneeded or out of date. This would require migration of data from one storage device to another. For instance, 1/2" magnetic tapes were used for digital data archives but since disappeared [23]. The form the data take also is a factor. For instance the programming language format of the data may not be popular enough in the marketplace to retain itself, though there are languages that have established themselves enough so there will be a place for them in the near future. The Objective-C language despite being fairly popular in 2013 has been on a steady decline since 2014 [15]. The technology to read and write the data will in time also be obsolete. Even if the data carrier itself survives, if we cannot read it then the data become a moot point. Today, how often does one encounter beta tape players? Do you know what Betamax is? Thus the library should be built on a device that would not require constant turnover but in a form that, with enough foresight, will last the longest amount of time.

2.2 Possible Solutions

To begin with a list of current possible answers to long term retention:

- Microfilm
- Magnetic Disks
- SSD (Solid State Drives)
- DNA Data Storage
- Optical Data Storage
- Holographic Data Storage

A possible solution is physical storage microfilm which can be held for up to 500 years with no notable degradation [23]. However given that the library can and will update technology as time goes on, it will require hardware that enables the user to add collected information. The information gathered through the library containing facts, figures, history, etc., may at times be corrected, thus the need to maintain a read and update capability. Currently, magnetic disks are replaced every 5-7 years and solid-state drives at about every 10 years [5]. DNA data storage has a half-life of 500 years in harsh atmospheric conditions like high temperatures, and in a study was theoretically able to hold information for 2000 years [5]. However accessibility would present an issue. Optical data storage devices can maintain stability and readability for 100-1000 years and holographic data storage is able to have a lifetime of up to 50 years [5].

Of the available options, optical storage and holographic storage are the higher tier choices. However at present, holographic storage devices are still relatively new and costly [18]. If the prices are not lowered current data storage industries may not invest enough for it to become mainstay. Without a significant enough base of users the costs would remain high and the need for tools to read it would be few and far between, and in relatively short time possibly fade out like beta tapes during the time of VHS cassettes.

Chapter 3

Requirements for the Library

3.1 Error-Correcting code

Redundancy is necessary to allow a strong level of availability and in the case of corrupt code, issues during transfer, and damage or loss of the storage device itself. Given the possible size of the storage system it is necessary to lower disk usage, improve the speed of writes when necessary though the system wouldn't have too much outside of initial start-up, but more importantly improve the usage of bandwidth for file sharing and redundancy. Bandwidth, which cloud-computing relies heavily on, would be the limiting factor. [26]

A basic form of redundancy is straight replication, one for one, which would have full copies of all data in alternative locations. It could be in the same system, network, etc. This would require significant storage space as any new addition would need to be replicated into any other backup storage locations.

Erasure coding is another another form. The main point is that the original object can be recreated from any m fragments and recoded into n fragments which are stored separately, [26], where the combined size of m fragments is approximately equal to the original object size. The goal is to increase redundancy and allow the same level of availability to be achieved with much smaller additional storage. In the tests run in [26], the authors concluded coding is not as useful for extremely high server availability due to the increased complexity introduced via erasure codes. Given the potentially large amount of information stored in the library, the server will need to be available almost all of the time. One example of erasure coding is the Reed-Solomon code that was standard up to the mid-1990's [4]. Reed-Solomon works well only in smaller scales which is why other faster and more scalable input size variants were created, such as Tornado, LT and Raptor codes. In the case of Reed-Solomon, c repair packets are generated and sent for every r data packets and the correct delivery of any r of the r+c packets transmitted is sufficient to reconstruct the original r data packets [4]. Dimakis et al. used a variation of erasure codes called Regenerating Codes to minimize the repair bandwidth. Lakshmi and others build upon Dimakis' Regenerating Codes to create minimum bandwidth regenerating codes to check the integrity of received data and correct the channel errors in the data iteratively [19]. Similarly, the form of the long term data retention will need erasure coding of its own to maintain integrity of the data while also not over-utilizing the bandwidth between storage devices.

In distributed storage systems, redundancy must be continually refreshed as nodes fail or leave the system; this involves large data transfers across the network [13], alternatively, there are erasure codes that can be repaired without communicating the whole data object. There is a trade-off between repair bandwidth and storage which regenerating coding can achieve optimally. It balances between minimum-storage regenerating codes, the maximum distance separable codes that can be efficiently repaired, and minimum-bandwidth regenerating codes, have minimum repair bandwidth.

Regenerating codes address the issue of some of erasure codes in distributed systems in that regenerating codes are not as complicated. In [13], the authors proposed regenerating codes could be potentially applied to archival databases. In distributed archival storage or backup, files are large and infrequently read thus error correcting codes may offer benefits in redundancy, reliability and repair bandwidth. Erasure coding like the Hybrid strategy complicates the architecture with minimum if any benefits and thus should be avoided. Ismail tested averages for bandwidth against replication, ideal erasure codes, hybrid, minimum-storage regenerating codes and minimum-bandwidth regenerating codes [17].

3.2 Integrity

While saving the data for future generations is the goal, knowing that the data is accurate is key. Data must maintain integrity, accuracy and consistency, and authenticity, the data originating from its source, for as long as the data exist; to guarantee authenticity the data requires a digital signature or an equivalent. However with computational power ever increasing, it is only a matter of time before attackers have the capabilities to break the digital signatures. Thus these signatures must be updated regularly and either become longer with each update or an alternative signature method needs to be used.

Data integrity is defined as data that is untouched after its creation. In the case of the library database, the data was scanned or retrieved from the origin source.

There cannot be any doubt that the data viewed is directly from the original. For example, if the Mona Lisa portrait is scanned into the database unfiltered, unaltered and completely one-for-one, the data would be authentic from the original. This also applies when users retrieve the data, it is from the source of the library database and not from some third party. It cannot be authentic if the connection from user to data source cannot be confirmed. Another key word is accuracy, as not a single value or portion can be changed or lost to keep accuracy. If the scanned copy of the Mona Lisa lost a pixel, or a pixel was changed, it is no longer accurate thus lacks integrity. Integrity must guarantee the source and the intention of the source.

The regularity of renewing time-stamps will need to be assessed. If the time-stamps occur too far apart then issues can occur without detection. The data could be manipulated, decayed, or lost over time. Too many time-stamps and the storage space could fill up unnecessarily with constant time-stamps and run the system excessively which can be costly in regards to time and materials depending on format. In the experiments for the long term storage system LINCO, the goal is to maintain integrity, authenticity and confidentiality of data. Though the confidentiality aspect is unnecessary for our purposes, in their 100 year test they found it optimal to renew timestamps every two years due to typical storage hardware maintenance service intervals [6].

3.3 Authenticity

The signature updating methods can be done through notarization or time-stamping. Notarization in this sense requires a third party who will update the signature. This third party would need to be trusted. If the third party altered the data or approved false data then the authenticity would be lost. Alternatively time-stamping also requires a third party but includes all previous signatures to be checked each time. This would obviously cause the file sizes to grow and take longer each time they are updated. Thus it becomes a matter of space and time. A keyword for authenticity is consistency, how the data remains consistent and delivered to users in the same way, so that the usability of data is not different from one user to another. Authenticity ties into integrity, for without integrity, authenticity is not possible.

In an article from 2014, "A Performance Analysis of Long-Term Archiving Techniques" several schemes are tested for their storage space and verification time over 95 predicted years [15]. The authors sampled four different schemes: Martín A. Gagliotti Vigil's Notarial Scheme, Content Integrity Service, Advanced Electronic Signatures, and Evidence Record Syntax. These were evaluated for their time spent on signature verifications, hashing, non-cryptographic operations and analyzed computational bottlenecks of the implementations and how key sizes affect performance in the long term [15]. The testing concluded that Notarial Scheme is the most efficient if the verifier checks few documents and accepts strong trust assumptions. Evidence Record Syntax is efficient for many documents if the verifier trusts time-stamp authorities and is the one of the four that performs better with large keys.

The information stored may also be out-of-date, rendering it either useless or only useful in a historical context. There was a time when it was commonly believed that the Earth was the center of the universe and that the sun revolved around it. Though factually wrong it could be useful to know the era societies in which this belief flourished to see how it may have affected their educationally, culturally, and/or in their literature. Thus the necessity to time-stamp the data upon entry into the library.

In a similar vein, data entered should be entered with relevant metadata, such as the original format or form it was created in and/or with, acquisition, and other information. In the article by Michael Day, et al., the British Library created a draft profile for data entry that required background, acquisition, preservation intent, acquisition format, issues or challenges, and profile metadata [12]. This provides a strong basis for data entered in the library to assist in authenticity. The background and acquisition format would be significant for content historical purposes. A time and date for data entry that will inform when the data was entered into the library. Metadata would need to be entered for ease of searching within the library database. The final form can be left to the library data center committee referenced in chapter 5.

Chapter 4

Disasters

4.1 Potential Disasters

There are a number of hurdles to overcome in order to maintain a globally free library. Establishing its creation or how the data is housed, the multitude of ways the data could be damaged or tampered with, and accessibility and rights to the data. A meteorite could come crashing down on the data center and destroy the library. The chance is small but it is there. Other natural disasters such as floods, earthquakes, and fires are less damaging but more frequent. In this chapter we will discuss the potential ways for disaster to strike the library, focusing more what is likely and split into two parts; accidental and purposeful.

4.2 Accidental

4.2.1 Naturally Caused

Natural damage is something that can be counted on occurring eventually yet may not be able to be predicted in a reasonable way. There are seasons for tornadoes to occur yet when they form may not be known until its too late. There are a number of ways nature can damage a data center:

- Meteorological (Tornado, hurricane, hail, thunderstorms)
- Geological (Earthquakes, sinkholes, landslides)
- Astronomical (Meteorites, solar flares)
- Floods (Accidental and human error)
- Fire (Accidental and human error)

Each can be responsible for a number of problems for the centers besides being a threat to the human element. Regardless of the original cause, it will amount to basically the same outcome: structural damage to the building may force it shut down for extended period of time if not indefinitely, destruction of the data storage devices themselves, and inaccessibility to the data. In each case it is a costly process as the price for repair and/or recovery can be significant.

In June, 2009, there a was a disruption of service at one of Amazon's data centers caused by lightning hitting one of its facilities and damaging its power supply to several racks of servers. The outage lasted for five hours and downed a number of cloud server instances [8]. A cloud instance is a virtual server instance from a cloud network. Since then Amazon updated its policies to compensate for failed instances by re-provisioning services across different availability zones [8]. In January 2010, a faulty routing device in another Amazon data center caused the service Heroku to be down for one hour. This had the same setup problem as the 2009 issue where the entirety of the service was in a single availability zone. Had the service spread out to other zones it wouldn’t have been a news worthy problem [7]. Besides the accessibility issue, there is also a monetary factor. In 2013, a 49 minute outage of Amazon.com cost the company more than $4 million in lost sales [30] as well as the possibility of customers going to alternative businesses due to the downtime. In 2011, the Tohoku earthquake and tsunami forced companies to file bankruptcy due to critical backup data loss in data centers [21]. The key points are disaster prevention and data availability .

To maintain reliability and accessibility in a cloud service, the library must be able to store the information in multiple locations throughout the world. It must rely on the possibility that the cloud service in one availability zone will fail, no matter how far reaching, and that can and will halt service for possibly an extended amount of time. The library cannot "put all its eggs in one basket" regardless of subject or division.

4.2.2 Human Error

Many disasters would be related directly to or would lead back to human error. It could range from failed upgrades or updates, bad coding or configurations, or not setting up the data center properly from the onset. The issues may not be prevalent from the onset but problems will arise in time. In April, 2011, Amazon lost service in their Elastic Block Store (EBS) for four days due to a network configuration change [29]. It was supposed to be a normal upgrade to the capacity of the primary network. However a shift to one of the routers did not execute properly causing the primary network to fail and the secondary network could not handle the traffic. This problem extended to the Relational Database Service which relied on the EBS for database and log storage. Even with redundancies in place the problem persisted for days.

4.3 Purposeful

4.3.1 Hacking

A database is no good if the data cannot be trusted. Entry into the library will be discusses in chapter 5, but what about the data housed within? Error-correcting code should catch flipped bits and the like, however. The database needs to have security to prevent viruses from entering and altering or deleting the information contained. In an article by Yingxin, a study was performed that showed in a space of six months within two years the amount of virus attacks increased significantly from 96754 to 371652 between five University Digital Libraries [32]. Of those viruses, trojans were the vast majority over worms, backdoor, and others. Hackers could gain access via their created viruses, unintended back door in operating systems or other software means. They can steal or lock the information away and hold it for ransom. Given the library will be publicly viewable, a hacker could prevent access to key information for studies or from new information being added. They have the possibility to disrupt service and cause system failures and could alter information in the library causing it to lose integrity. It is imperative that the software of the long term data retention library be secure and updated frequently. The library will require an anti-virus software either of its own design or bought commercially like Trend Micro or McAfee.

4.3.2 Legality

In a perfect world, information would be accessible to any and everyone who requests it. However that is not one that exists currently and there are laws that must be obeyed. A library may be voluntarily given items by the creators or an institute can have the right to hold a copy of works by their employees.

Creators are giving the library the rights to use, display and share their works under certain, if any, restrictions. Libraries may pass books between each other if they establish the proper permissions. This is relatively easier to handle with physical items than with digital media. In the latter case once something is accessible online publicly, it is near impossible to retract. Thus any item housed in the library data center would need explicit permissions to share the entirety of the work before accepting it. An institution willing to share its collection would be amazingly substantial, but it would require all items to not infringe on copyright. In an article by J.M. Ashworth, the British Library has the right to obtain everything published in the United Kingdom [3]. However the items stored are not available for lending. In the same article it mentions the British Library made a deal with IBM to utilize a Digital Library System "to preserve and access electronic materials indefinitely even if the formats in which they originated are long dead" [3]. This was extended further in April 2013 to include select non-print content. Other European national libraries proposed a pan-European distributed digital library based on national collections and accessible to European citizens. A similar proposition would need to be utilized for a grand scale to ensure legality of all items with the library data center.

The legality of the permission must extend beyond an individual nation as well.

In article [14], the University of Michigan at Ann Arbor was willing to grant access to its 7,800,000 book collection to Google Library which included copyrighted materials. Rutgers library is currently working with Google by scanning items however no one has access to the full text of copyrighted works. A Chinese author brought a lawsuit against Google for infringement of their work and require compensation [14]. The authors explain that according to Article 22 in China's Copyright Law as long as the author's name and work's detail are pointed out and no infringement occurs then they can use the copyrighted work without contacting the author and obtaining their permission. In the case of China-based data, it would be best to maintain permissions by consulting the China Written Works Copyright Society (CWWCS) for related media. Similar agreements would be required with other nations and like Google Library an opt-in policy will need to be established.

4.3.3 War

Worst case scenario of a human-caused incident is large scale attacks. Whether by a terrorist group or war, one of the greatest fears is of a nuclear attack. In a simulation conducted of a 10-kiloton nuclear explosion in a large populated U.S. city, the authors review the effects of the damage and electromagnetic pulse (EMP) that generate thereafter on an electrical power system [25]. The authors list three types of EMP components that are generated by a nuclear blast. The first, E1, disrupts and damages electronic-based control systems including computers and protective systems. The second component is less significant to electronic systems and does not need to be elaborated. The third, E3, is as they describe as a "longer-duration pulse that creates disruptive currents in long electrical transmission lines, resulting in damage to connected electrical supply and distribution systems." [25]

Also listed within the article are major impacts of EMP components on electric power systems within a three mile radius:

- Electronic circuits and programmable logic controllers embedded in most electronic control and telemetry equipments are destroyed.
- High-amplitude surge currents, similar to those caused by geomagnetically induced currents from solar storms of equivalent intensity, are induced in transmission lines.
- The substation equipment located at the terminals of EMP-affected transmission lines is put at risk of permanent damage because of high current and voltage surges.
- Excessive currents are induced in transformer cores, resulting in overheating and fires.
- Transformer saturation and direct-current offsets tend to add inductive load to the system and cause voltage collapse, as available volt-ampere reactive (VAR) compensators are stretched beyond their capacity.
- Relays and control equipment malfunction and trigger unintended or undesirable system actions or switching.

All told, the effects on the power system are catastrophic. The impact area and surroundings would take many years to recover. Even if the storage area or data center were to be shielded from the blast, the electrical grid would be worthless and any on-site generators would most likely be damaged as well. Even if the storage was underground, the shock and EMP would still render the site almost worthless. If the data was stored underground, the best case scenario would be the data storage would be undamaged. However without a source of power the systems sustaining it won't last and adding or retrieving the data would be difficult if not impossible.

4.4 Disaster Solutions

Tor the sake of redundancy cloud servers or databases should be stored in no fewer than three locations. The Digital Library System of the British Library has four storage nodes contained in London, Boston Spa, Edinburgh and Aberystwyth [12]. For the sake of worst-case scenarios with drastic weather conditions and war, the storage locations should be located far from each other and exist on separate continents. For the Digital Library System of the British Library, if a catastrophe were to befall any of the four cities the other three are far enough away to be unaffected by most. However in a truly cataclysmic event like war, the safety of all four nodes is put into question. The location must be in a fairly neutral area, not in an area known for conflicts or severe weather conditions of any variety. Though unusual weather patterns may occur, it is possible to ascertain such locations through history, analyzing current patterns and forecasting years into the future.

Accessibility to the information needs to be protected. The paths to the data need to be secured and reliable. If there is a severed connection anywhere, there should be multiple paths leading back to it. Relying only on a single path or just two, main and a backup, will eventually lead to disaster. The alternative lines should be link-disjoint from the main so as to avoid a disaster that can take out multiple lines as purposed by [16]. The authors believed, and common sense dictates, that the loss of a destination node should not halt anyone's access. The authors worked towards an integrated Integer Linear Program to find a solution simultaneously for "content placement, routing, and protection of paths and content" [16]. The authors discovered that using an interlaced Integer Linear Program formulation for content placing and routing reached the best-case scenario for network optimization and data center resources in a real-life scenario and ran in weakly-polynomial time as it depends on the number and size of the input. These backup paths can also be used in case the bandwidth demands increase enough to cause slowdown or disconnection. They can be used to redirect the demands and alleviate the network.

The location should also be prime for resources. It should not be in a location with lots of interference. It will need adequate space for possible expansion. In the article by Xaiole Li et al., they proposed a few key points [21]. First is to assess the risk analysis on location if a disaster occurs and the potential cost lost if one were to occur. Alternatively it should also value when a disaster occurs how best to save the data via distribution to backup servers far enough away from the disaster's influence. This relies on the network's capabilities and the speed it takes to begin the backup operation and how long it will take to move said data before the disaster destroys the data or causes a loss of power. The use of evacuation latency is to evaluate the time it takes to move the data off site. However depending on what the disaster is and how much it may repeatedly occur, the algorithm presented in [21] shows there is some variance in the optimal solutions as it needs to find abide by given backup objectives to address the particular disaster. For instance if a massive tornado hits one of the primary data centers, the backup locations should not reside in a nearby city. Though it would be beneficial by proximity thus able to transfer data faster in an emergency, the disaster could hit the backup location as well. If no more optimal backup solutions exist, then there could be significant data loss. The algorithm of [21] finds optimal backup locations by calculating the risk factor that a backup center holds and compares to the others. Thus the main points are disaster risk distribution and evacuation latency to find optimal locations that provide low risk and multiple back-up solutions.

Chapter 5

Data Review

5.1 Filtering

Accepting all data would be a reckless endeavor. If the information already exists in the database, or within a journal that already exists, then it may create duplicates, taking up precious space in the data centers. The library should not allow everything that exists to be entered. The goal is not to be another Wikipedia where data can be altered by anyone at anytime. Though the idea is sound, data can be entered with no regards to accuracy or proof. Users may change existing data due to their inaccurate information or as a prank. Without a proper filtering process the data can be arbitrary. There also needs to be a proper entry system so that the data is entered correctly and properly labeled. As mentioned in chapter 3, there needs to be a framework profile for the entries to maintain authenticity. Thus there needs to be a selection process of certified individuals who can consistently and accurately add new data to the library.

5.2 Committee

My suggestion is to use a committee to maintain and update the data depository, such as the World Data System (WDS) [27]. An offshoot of the ICS (International Science Council), the World Data System's mission is to promote long-term stewardship of and universal and equitable access to quality-assured scientific data and data services, products and information across all disciplines in the Natural and Social Sciences, and the Humanities. WDS's strategic targets perfectly exemplify the tasks necessary for this long term data retention plan. The explore long term data stewardship, compliance to agreed-upon data standards and conventions, and improving sustainability, trust and quality of the data.

Though the World Data System is relatively new, established in 2009, it grows from the World Data Centres and Federation of Astronomical and Geophysical Data Analysis Services originated in 1957-1958. WDS was created as a replacement when the previous group was not able to react to the modern data needs of the growing world. In WDS' Strategic Plan in 2019, one of their key goals is to ensure the integrity of science, and the long-term preservation of data underlying scientific knowledge from all disciplines [28]. Though the focus is on scientific datasets, their goal can easily be extended to data retention of all varieties.

Data should be approved through a group of scientists, scholars, historians and librarians. Only relevant "important" data should be selected for long term preservation: Mathematical and scientific equations, history of various cities, states and countries, star charts, significant works of art and literature. It should be comprised of information, scientific theories and datasets. The information stored would be like what we would send in the Voyager Golden Record, only greatly expanded upon. All information stored within should be freeware, with no need of concern about timing or licenses.

Chapter 6

Data Centers

6.1 Data Center Locations

To house a library in multiple locations one needs to address most if not all of the issues brought up in previous chapters. Natural disasters, energy efficiency, human element all play a part of this endeavor. As of June 2018, the largest data center in the US is housed in Las Vegas, Nevada [1]. Titled the Switch SuperNAP, it occupies 3.5 million square feet. It is adjacent to the Tesla Gigafactory and is completely run by renewable energy with 100% Clean Energy Index. It is also protected from some natural disasters and possesses security as it is a part of the Tahoe Reno Industrial Center which is surrounded by a solid concrete wall 20 feet high.

A new larger data center named Kolos in Ballangen, Norway is in the planning stage as the largest center, [1]. It is an opportune place as the location near the Arctic will lower energy cost by 60% due to cold climate and access to hydroelectricity. The Kolos facility is surrounded by hills and water which grants it protection from most physical disasters. In the book by Jakob Christensen, the Nordic region is a global leader in the digital economy [10]. In all of Europe it possesses Long-Term Evolution infrastructure and high penetration of fibre broadband services domestically and internationally as well as the highest percentage of Nordic enterprises using cloud computing services. Through data collected on the important factors for data center installation in general, what the Nordic region would provide is a reliable power supply, low energy prices, political stability, time-to-market, abundance of energy and other resources and scalability. The region is also connected to the UK, Europe and the U.S. with more connections planned including a polar route from China.

A third possible library is in Singapore, a location placed top of Asia-Pacific regions in 2017 [20]. Singapore had two new data centers created between 2015-2017 and cloud operators had become the largest occupier for data centers. The capacity in Singapore reached 370 megawatts with another 100 more in a future expansion. Google built one of its data centers in 2013 due to its fast growing Internet market, reliable infrastructure, and transparent and business friendly regulations [9]. The location is also one of the lowest to be hit by natural disasters of ASEAN, Association of Southeast Asian Nations [17]. Not only does it possess one of the lowest climate risks from 1998 to 2017 but there was no recorded deaths or significant losses in the entirety of 2017. It makes a prime location for one of our data centers.

6.2 Data Center Architecture

The data center itself needs to be maintained which will require a scheduler to allocate workloads to resources. In an article by Georgios Andreadis, he and his team developed a reference architecture for data center scheduling and use it to analyze academic and industry-designed schedulers for general use [2]. The basic proposed model supports a diverse set of scheduling operations in data centers that appear in peer-reviewed research and in practice. The scheduler needs to follow the policies given for the data center besides the established ones, like with Shortest-Remaining-Time-First which alters the priority on tasks that are shorter to be higher than others.

The referenced architecture is designed with validity and usefulness in mind for use in current and possible future datacenter scheduling and with a real-world purpose. The key principles of the scheduler design are labeled as follows:

- Components with Clearly Distinct Responsibilities
- Grouping of Related Components
- Separation of Mechanism from Policy
- Scheduling as Complex Workflow, Matching the System Model
- Hierarchical Scheduler with Shared Control

Each point is divided up further within the related components: the four major tasks are job processing, task processing, scheduler management and resource management. In their results the authors found a number of key differences between industry and academia, particularly in regards to computation correctness, consistent state of the system, authorization issues, job cleanup and job completion. The academic reference scheduler will be useful in design of the library data center architecture. It can be utilized to lay the ground work, comparing existing models prior to the setup phase and develop its own scheduler to maintain the proposed data center. Georgios Andreadis and his team plan to expand its capabilities and support on a global scale, thus this architecture will be more of an asset in the future [2].

6.3 Data Center Cost

The creation or renting of data centers also comes with a cost for the electricity and bandwidth. Ying Zhang et al. offer a solution in their research on how to minimize the cost of the two factors [33]. There is not a way to get 100% data in the market and thus there is a level of uncertainty. There needs to be a knowledge of the workload and how much electricity it will require to run optimally. Ying Zhang et al. note that the electricity cost is from the day-ahead market, in essence the center would bid for what they require and suppliers will offer an amount and price [33]. The bidding will have to go back and forth until a settlement is reached. Since this is paid ahead, there are a few other stipulations that need to be addressed.

[33] list the following:

- In case that the day-ahead committed supply matches exactly the actual demand, there is no real-time cost.
- In case of under-supply, (i.e., the committed supply is less than the real-time demand), the cloud service provider will pay for extra supply at the real-time price.
- In case of over-supply, the system needs to reduce the power generation output or pay to schedule elastic load to balance the supply, both incurring operational overhead and consequently economic loss. In this case, the cloud service provider will receive a rebate for the unused electricity.

A key point to factor is the global load balancing, the ability to move the data traffic through multiple data centers in different locations to maximize security and efficiency. The data center would have to optimize the global load balance and bidding to keep the operating costs at a reasonable level. Thus the key points that must be factored together are workload, electricity, geographical load balancing, and bandwidth cost. With limitations the author's algorithm is able to reduce the cost up to 20%. The limitations pertain to assumptions during the algorithm's creation, which is if the day-ahead and real-time markets are exclusive, that the cloud service provider has negligible market power and the objective function's gradient cannot be computed explicitly. The former would be discovered during the bidding, the market power does not factor for the library data center as it would not be local only and the objective function's gradient would need to be solved. In the proceedings from Hong Xu and Baochun Li, they take a different step and include the calculations for bandwidth demand from the data centers [31]. The authors utilize the requests at different locations to save on cost by paying for the higher used nodes during peak times and reducing the cost from nodes with fewer requests. Both algorithms are a step in the right direction in calculating costs of running the library data centers.

Chapter 7

Cloud Computing

7.1 Why Cloud Computing?

One of the reasons for choosing cloud computing is the high probability of sustaining over time. Jose Picado, Willis Lang, and Edward Thayer discuss how long a database will survive before being dropped by the users [24]. Survivability is tied to revenue, usage, upkeep, and the like. The study focused on Microsoft Azure SQL Database in three separate regions and used a relative microcosm to predict whether a database will survive beyond 30 days. They tested over five months in a single year and divided over three regions and acquired data on duration from creation to cancellation over three different database editions [24]. The editions were Basic, Standard, both used on remote storage, and Premium which used local storage. All editions vary in multiple service level objectives like redundancy and back-up retention. For the authors' sample, the Premium's database size was significantly smaller than the others. They utilize random forests due to their ability to be trained for quick and accurate predictions and survival analysis tools that analyze the duration until an event occurs. In this case users dropping the database. One of the major observations discovered by the authors is "proportionally few databases in Basic and Standard edition switch to a new edition during their lifetime" [**Referene26**]. Using their results, they were able to obtain an accuracy of more than 0.80. The specifics of the Basic edition were able to successfully predict 92% of the databases that live longer than 30 days. Picado et al. used the tool to help providers know why databases are dropped and improve their policies and infrastructure [24]. Using their method, they could make confident predictions on average in 63% of Basic databases, 90.3% of Standard databases, and 71% of Premium databases. In all cases, the probability of sustaining databases is decent to well above average.

The second reason is the prevalence and growing market of cloud computing and in particular cloud services. The market for cloud services continues to grow with a steady climb. In a 2017 article by Muelen and Petty, Sid Nag, research director of Gartner, predicted by 2020 that all cloud adoption strategies will influence more than 50 percent of IT outsourcing deals [22]. The reason being, the author states, is that buyers will investigate cloud options first due to time-to-value impact via speed of implementation. The software as a service, specifically in this case cloud application services, revenue was much higher than expected for 2016 and they foresaw the trend to continue for years. Sid predicted by 2018 the software as a service to reach $ 71.2 billions of U.S. dollars however in the forecast report created in November of 2019, it had actually reached $ 85.7 billion [11]. In the 2017 article Sid stated that as of 2016 "approximately 17% of the total market revenue for infrastructure, middleware, application and business process services had shifted to cloud" [22]. In Sid's article from 2019 he stated that "by 2022, up to 60% of organizations will use an external service provider's cloud managed service offering, which is double the percentage of organizations from 2018" [11]. According the Sid Nag, cloud computing will continue to grow and be ingrained in various ways.

For these reasons the long-term data retention method should utilize cloud databases from the start. Given the consistency of cloud databases and how they they are forecast to not only be sustained for years but will continue to grow for years to come in personal and business fields.

Chapter 8

Conclusion

There is a need for accurate data for current and future generations to utilize. This can be done by creating a long term data retention library. With the current technology, cloud computing has shown to be a growing database structure and used for data storage for many companies. The retention library must be able to maintain authenticity via a committee to approve of new entries and security software to keep the integrity. It will need to be housed in three nodes at distant locations that are less likely to be damaged from natural disasters. In the event a disaster occurs, it must be able to continue running if a data center is incapacitated. It would have been beneficial to investigate the British Library and the Library of Congress for current data on their proceedings, how they deal with problems and the costs for maintenance. In the future, holographic data storage could be a suitable option if the cost lowers and the read/write speed catches up or surpasses standard cloud computing. Regardless of the case, we need to find a suitable storage device to house our data lest it be lost in time or our machinations.

Bibliography

[1] M. Allen. *And The Title of The Largest Data Center in the World and Largest Data Center in US Goes To...* 2018. URL: https://www.datacenters.com/news/and-the-title-of-the-largest-data-center-in-the-world-and-largest-data-center-in (visited on 01/05/2020).

[2] G. Andreadis et al. "A Reference Architecture for Datacenter Scheduling: Design, Validation, and Experiments". In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. SC '18. Dallas, Texas: IEEE Press, 2018, 37:1–37:15. DOI: 10.1109/SC.2018.00040. URL: https://doi-org.proxy.libraries.rutgers.edu/10.1109/SC.2018.00040.

[3] J. M. Ashworth. "British Library's role as the national library and its commitment to electronic international library development". In: *Proceedings 2000 Kyoto International Conference on Digital Libraries: Research and Practice*. 2000, pp. 12–15. DOI: 10.1109/DLRP.2000.942147.

[4] M. Balakrishnan et al. "Maelstrom: Transparent Error Correction for Communication Between Data Centers". In: *IEEE/ACM Transactions on Networking* 19 (2011), pp. 617 –629. ISSN: 1558-2566. DOI: 10.1109/TNET.2011.2144616.

[5] W. Bhat. "Bridging data-capacity gap in big data storage". In: *Future Generation Computer Systems* 87 (2018), pp. 538 –548. ISSN: 0167-739X. DOI: https://doi.org/10.1016/j.future.2017.12.066.

[6] J. Braun et al. "LINCOS: A Storage System Providing Long-Term Integrity, Authenticity, and Confidentiality". In: *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*. ASIA CCS '17. Abu Dhabi, United Arab Emirates: ACM, 2017, pp. 461–468. ISBN: 978-1-4503-4944-4. DOI: 10.1145/3052973.3053043.

[7] C. Brooks. *Heroku learns from Amazon EC2 outage*. 2010. URL: https://searchcloudcomputing.techtarget.com/news/1378426/Heroku-learns-from-Amazon-EC2-outage (visited on 01/05/2020).

[8] C. Brooks. *Users undeterred by Amazon EC2 lightning snafu*. 2009. URL: https://searchcloudcomputing.techtarget.com/news/1359572/Users-undeterred-by-Amazon-EC2-lightning-snafu (visited on 01/05/2020).

[9] Google Data Centers. *Our first data center in Southeast Asia*. URL: https://www.google.com/about/datacenters/locations/singapore/ (visited on 01/05/2020).

[10] J. Christensen et al. *DATA CENTRE OPPORTUNITIES IN THE NORDICS*. Nordic Council of Ministers, 2018, p. 7. URL: https://norden.diva-portal.org/smash/get/diva2:1263485/FULLTEXT02.pdf (visited on 01/05/2020).

[11] K. Costello and M. Rimol. *Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17in 20207*. 2019. URL: https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2020 (visited on 01/05/2020).

[12] M. Day et al. "Implementing Digital Preservation Strategy: Developing content collection profiles at the British Library". In: *IEEE/ACM Joint Conference on Digital Libraries*. 2014, pp. 21–24. DOI: 10.1109/JCDL.2014.6970145.

[13] A. Dimakis et al. "Network Coding for Distributed Storage Systems". In: *IEEE Transactions on Information Theory* 56 (2010), pp. 4539 –4551. ISSN: 1557-9654. DOI: 10.1109/TIT.2010.2054295.

[14] Y. Guo, Y. Liu, and Z. Yu. ""Google Library": Some Copyright Infringement Concerns in China". In: (2010), pp. 2053–2056. ISSN: null. DOI: 10.1109/ICEE.2010.519.

[15] J. H. *Worst Programming Languages to Learn in 2018*. 2018. URL: https://www.codementor.io/blog/worst-languages-to-learn-3phycr98zk (visited on 01/05/2020).

[16] M. F. Habib et al. "Design of Disaster-Resilient Optical Datacenter Networks". In: *Journal of Lightwave Technology* 30.16 (2012), pp. 2563–2573. ISSN: 1558-2213. DOI: 10.1109/JLT.2012.2201696.

[17] M. Ismail. *ASEAN nations hard hit by natural disasters*. 2018. URL: https://theaseanpost.com/article/asean-nations-hard-hit-natural-disasters (visited on 01/05/2020).

[18] Vaniya Vimal Lakhani Amitkumar and Gurnani Honey R. "Holographic Data Storage Technology: The Future of Data Storage". In: *international journal*

on recent and innovation trends in computing 6.1 (2016), pp. 111–114. URL: https://ijritcc.org/index.php/ijritcc/article/view/1389.

[19] V. S. Lakshmi and P. P. Deepthi. "Error Correction Scheme for Regenerating Code Based Distributed Storage Systems". In: *Proceedings of the 2018 International Conference on Communication Engineering and Technology*. ICCET '18. Singapore, Singapore: ACM, 2018, pp. 5–7. ISBN: 978-1-4503-6454-6. DOI: 10.1145/3194244.3194246.

[20] C. Li. *Data Center Investment: A Rare Opportunity For the Right Investor.* 2017. URL: http://www.cushmanwakefield.sg/en-gb/research-and-insight/2017/data-centres-report (visited on 01/05/2020).

[21] X. Li et al. "Disaster-and-Evacuation-Aware Backup Datacenter Placement Based on Multi-Objective Optimization". In: *IEEE Access* 7 (2019), pp. 48196–48208.

[22] R. Muelen and G. Pettey. *Gartner Forecasts Worldwide Public Cloud Services Revenue to Reach $260 Billion in 2017*. 2017. URL: https://www.gartner.com/en/newsroom/press-releases/2017-10-12-gartner-forecasts-worldwide-public-cloud-services-revenue-to-reach-260-billionin-2017 (visited on 01/05/2020).

[23] F. Müller et al. "PEVIAR: Digital Originals". In: *J. Comput. Cult. Herit.* 3.1 (July 2010), 2:1–2:12. ISSN: 1556-4673. DOI: 10.1145/1805961.1805963. URL: http://doi.acm.org.proxy.libraries.rutgers.edu/10.1145/1805961.1805963.

[24] J. Picado, W. Lang, and E. Thayer. "Survivability of Cloud Databases - Factors and Prediction". In: *Proceedings of the 2018 International Conference on Management of Data*. SIGMOD '18. Houston, TX, USA: ACM, 2018, pp. 811–823. ISBN: 978-1-4503-4703-7. DOI: 10.1145/3183713.3190651.

[25] E. C. Portante et al. "Simulating the potential impacts of a 10-kiloton nuclear explosion on an electric power system serving a major city". In: *2013 Winter Simulations Conference (WSC)*. 2013, pp. 2508–2519. DOI: 10.1109/WSC.2013.6721624.

[26] R. Rodrigues and B. Liskov. "High Availability in DHTs: Erasure Coding vs. Replication". In: (2005), pp. 226–239.

[27] World Data System. *International Council for Science*. 2019. URL: https://www.icsu-wds.org/ (visited on 01/05/2020).

[28] World Data System. *WDS Strategic Plan 2019-2023*. 2019. URL: https://www.icsu-wds.org/publications/strategic-plans (visited on 01/05/2020).

[29] The AWS Team. *Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region*. 2011. URL: https://aws.amazon.com/message/65648/ (visited on 01/05/2020).

[30] S. Withers. *The Cost of a Cloud Outage*. 2014. URL: https://www.business2community.com/Cloud-computing/cost-Cloud-outage-0925420#b4pjXoK2JIgOUkHk.97 (visited on 01/05/2020).

[31] H. Xu and B. Li. "Cost efficient datacenter selection for cloud services". In: *2012 1st IEEE International Conference on Communications in China (ICCC)*. 2012, pp. 51–56. DOI: 10.1109/ICCChina.2012.6356938.

[32] Y. Zhai and X. Liu. "Research on network security measures of digital library in university". In: (2011), pp. 3630–3633. ISSN: null. DOI: 10.1109/ICMT.2011.6002190.

[33] Ying Zhang et al. "Joint Bidding and Geographical Load Balancing for Datacenters: Is Uncertainty a Blessing or a Curse?" In: *IEEE/ACM Trans. Netw.* 26.3 (June 2018), pp. 1049–1062. ISSN: 1063-6692. DOI: 10.1109/TNET.2018.2817525. URL: https://doi-org.proxy.libraries.rutgers.edu/10.1109/TNET.2018.2817525.