Plasmon White Paper

Archival Storage Total Cost of Ownership Analysis

A detailed financial analysis comparing the use of magnetic disk, magnetic tape and optical technologies within an archival storage environment

The purpose of this Total Cost of Ownership (TCO) analysis is to compare and contrast the quantifiable acquisition and operating costs for a cross-section of different storage technologies used for archiving data.

Author: Plasmon
Version: 4.1
Datum: 22 August 2007
Status: Release
Distribution: Released

Table of Contents

Table of Figures

1 Executive Summary

In the present climate of increasing regulations and the risk of crippling litigation, it is important for organizations to have secure, long-term access to valuable information. In developing a data archive strategy there are a range of archive technologies to choose from, each with its own strengths, weaknesses and costs.

The purpose of this Total Cost of Ownership (TCO) analysis is to compare and contrast the quantifiable acquisition and operating costs for a cross-section of different archive solutions. In order to develop a representative model for the analysis, an actual case scenario was used. The requirement was for a 40TB archive measured over 3 years of operation. The archive products selected for comparison are a Quantum i500 tape library (LTO3), a Plasmon UDO Archive Appliance (UDO2), a NetApp FAS 3020 (SATA disk), and two configurations of an EMC Centera (SATA disk).

In order to avoid subjective interpretation of the TCO figures, only clearly quantifiable costs were included in the analysis. All costs used were list price US$ values and were proportionally adjusted to fairly compare similar 40TB configurations. The costs included in the analysis were hardware and software acquisition, media acquisition, hardware and software maintenance, floor space and the cost for power and cooling.

Figure 1 below provides a high level summary of the results of the TCO analysis. The body of the report documents the methods used in gathering these results and provides a detailed analysis of the final figures.

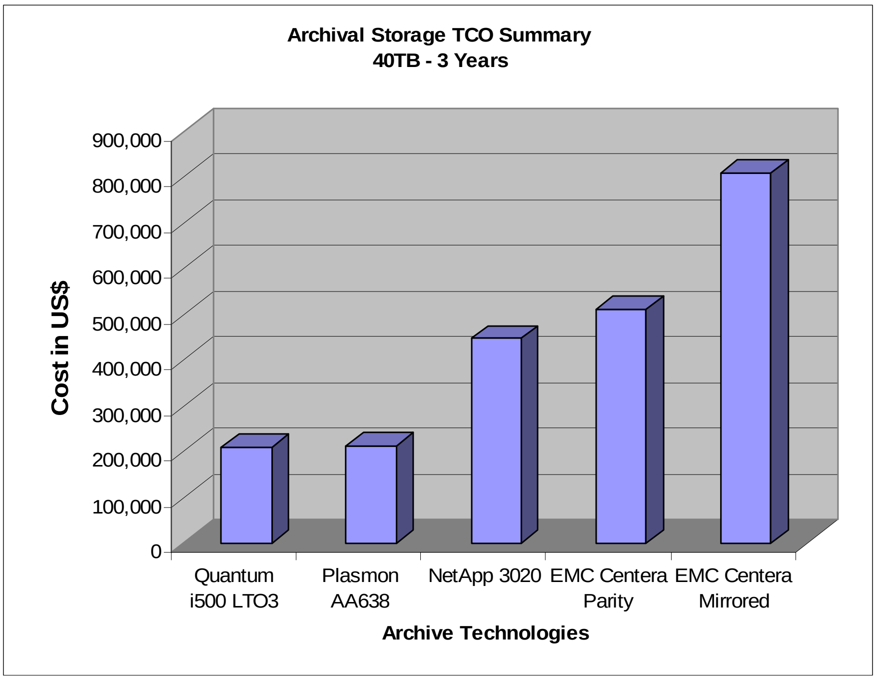

Figure 1 – Archival Storage TCO Summary

The results clearly show that the Plasmon UDO Archive Appliance and Quantum i500 achieve the lowest Total Cost of Ownership. Both the NetApp and EMC magnetic disk based systems were substantially more expensive than the Plasmon and Quantum solutions. In particular, the EMC Centera Mirrored configuration demonstrated a dramatically higher TCO than any of the other products.
The numbers gathered in this analysis do not account for all possible system and operating expenses, but do provide an accurate relative percentage for the selected technologies. The analysis strives to provide a factual starting point for readers looking to analyze the TCO of an archival storage environment and also touches on some of the additional business and technical considerations important in developing a successful and cost effective archive strategy.

2 Introduction

The demand for long-term record archives has increased sharply in recent years. Growth has been fuelled by new government regulations affecting data retention, the need to defend more effectively against litigation and the competitive imperative to maximize the value of organizational assets. These drivers have influenced large and small businesses across virtually every industry, as well as government and non-profit institutions.

Many organizations are now faced with the need to cost-effectively store and quickly retrieve large volumes of electronic content for years or decades. In many cases, the requirements of a long-term data archive have introduced completely new technical and operational considerations that can have serious consequences if not properly managed.

One area that is frequently discussed, but seldom properly quantified is the Total Cost of Ownership (TCO) for the acquisition, maintenance and operation of a digital archive. The purpose of this report is to identify and assess the most significant financial factors in a long-term record archive using the most common archival storage technologies: magnetic disk, magnetic tape and UDO. This analysis does not attempt to capture all possible costs, but to identify the most significant quantifiable costs in order to provide an even-handed TCO comparison among the different technologies.

TCO is a very important consideration when selecting a technology for use in an archival storage strategy, but it is by no means the only determinant. While this report does touch on additional strategic considerations, the primary focus is to analyze the financial impact of different archive solutions in the context of a long-term data archive.

3 Archive Case Scenario

In order to provide a realistic TCO assessment, the analysis uses the actual archival storage requirements from the office of a global financial institution. The organization has data archive requirements for the storage of stock trading transactions and customer financial records. The need to retain these records is driven by both internal corporate policies and governmental regulations that mandate the long-term retention of financial records.

The institution maintains their active records on protected, high performance, magnetic disk for 30 days before moving the records to an archive. Record access after 30 days is sufficiently infrequent as to warrant a more cost effective archival storage strategy that complies with their legal obligation for record retention.

Since this organization trades U.S. securities, one of the primary areas of concern is compliance with the Securities and Exchange Commission (SEC). SEC regulation 17a-4[1] governing broker-dealer transactions require record retention for up to 7 years and states that the archive media must use a “non-erasable, non-rewritable” format. This requirement for a non-alterable storage media is a critical component in establishing long-term data authenticity and appears frequently in archival storage regulations worldwide. As a result, the company requires an archive storage strategy that will meet very specific longevity and authenticity demands.

A total archive capacity of 40TB has been set as a target. With daily archive volumes averaging 8GB to 10 GB and factoring in yearly growth rates of 30-40%, the 40TB capacity is designed to meet their archive requirements for the next 3-5 years. They have thousands of networked users with a need to view historic records and have an average of 2,500 requests for archive data in a standard 8-hour working day.

[1] The SEC website is www.sec.gov. A full text of the SEC 17a-4 regulations can be found on www.law.uc.edu.

This particular scenario was selected for the analysis since it provides significant archive volumes, definable access patterns and very specific retention periods as required by industry regulations. While different industry regulations and corporate policies can vary dramatically, data retention requirements are becoming ever more common for many organizations and this particular scenario represents a "typical" case study that is well suited for the analysis.

Figure 2 summarizes the archive attributes upon which the TCO analysis is based.

Case Scenario Requirements	
Required Archive Capacity	40TB
Typical Record Retention Period	7-10 years
Average Daily Archive Volume	8GB to 10GB
Total Number of Records	> 6 million
Average Daily Archive Read Requests	2,500/day

Figure 2 – Case Scenario Requirements

4 Selected Archive Solutions

The TCO analysis compares the most common archive storage technologies from industry leading manufacturers including: magnetic tape, magnetic disk and UDO storage technologies. All of the selected products are marketed by their manufacturers as long-term archival storage solutions. The analysis compares the latest generation of products available at the time this report was written.

4.1 Magnetic Tape Archive

LTO3 tape was selected as the tape archive technology since it is one of the most popular tape formats in use today. LTO is a moderately priced professional class product with high capacity media. LTO tapes are also available in a WORM (Write Once Read Many) emulation format. When a WORM tape is loaded in an LTO drive, the drive will not allow previously written data to be erased and rewritten. WORM media is an important requirement for archive environments where data authenticity is critical.

LTO3 has an uncompressed media capacity of 400GB and a maximum compressed capacity of 800GB. The amount of achievable drive-based compression is very dependant on the size and type of files. Compression often provides little or no benefit for small and medium sized business records that may already be saved using some form of compression. With this in mind, the uncompressed 400GB media capacity has been used in this analysis.

4.2 Magnetic Disk Archive

With the continual decrease in the cost of magnetic disks, a number of specially designed magnetic disk archive products have entered the market. This analysis has chosen NetApp FAS 3020 and EMC's Centera (Generation 4 LP) products since they are targeted at the archival storage space. Both NetApp and EMC offer specific features to meet archive requirements including WORM emulation. Similar to WORM emulation on tape, NetApp and EMC use a software interface to prevent previously written data from being erased and rewritten. WORM capability is a key feature for this scenario and for many archive environments.

It is important to note that this analysis has selected magnetic disk solutions specifically designed for archival storage. While there are many "low cost" RAID systems on the market, most do not offer the management and authentication features required in a long-term archive so would not offer an appropriate comparison.

4.3 UDO Archive

Plasmons' UDO Archive Appliance (AA638) was selected as an alternative to magnetic disk and tape recording technologies. The UDO Archive Appliance is a network attached archive solution that uses a small amount of magnetic disk cache (2TB in this example) for simple connectivity and fast performance, with data securely archived on second generation 60GB UDO2. In this analysis UDO2 Write Once media has been selected. This true WORM (Write

Once Read Many) media cannot be physically altered once written providing the very highest standard for data authenticity, which is critical to the security of the data in this scenario.

5 Archive Configurations

In order to meet the required 40TB archive capacity complete archive systems have been configured using tape, disk and UDO technology., Figure 3 below details the specific configurations selected and the total usable capacity for each of the three technologies.

Media Type	Vendor	Product	Drive Count[2]	Media/Drive Capacity	Media Count	Usable System Capacity
LTO3	Quantum	Scalar i500	10	400GB	128	51.2TB
UDO2	Plasmon	AA638	6	60GB	638	38.3TB
Mag. Disk	NetApp	FAS 3020	70	750GB	-	43.5TB
Mag. Disk	EMC	Centera Parity	80	750GB	-	42.5TB
Mag. Disk	EMC	Centera Mirrored	128	750GB	-	41.4TB

Figure 3 – Selected Archive Configurations

Figure 3 also illustrates that it was difficult to select configurations that precisely match the 40TB archive capacity. Instead, the analysis has chosen fully populated systems that most closely match the 40TB requirement and mathematical adjustments were made in the financial model to compensate for the difference in total archive capacity. Refer to the Capacity Adjustment Calculation Section 6.1 for a complete explanation.

It should be noted that the LTO3 configuration required the selection of a larger capacity library than needed in order to accommodate a higher number of tape drives required for an archive. Section 5.1 provides a detailed explanation of the drive count calculation.

It should also be noted that the Centera product has been listed with two possible disk configurations: Parity (Content Protection Parity or CPP) and Mirrored (Content Protection Mirroring or CPM). These two architectures afford different levels of data redundancy with a trade-off in system capacity. The report authors understand that the most commonly recommended configuration is the mirrored version, but in order to avoid any incorrect assumptions, both configurations have been listed.

5.1 Drive Count for Tape and UDO

In order to satisfy 2,500 access requests each working day, the archive must have sufficient read bandwidth. Additionally, there must also be sufficient write bandwidth to accommodate the 8GB to 10GB of daily archive volume. In the case of tape and UDO, the ability to service these requests is dependent on the number of available drives and their performance. It has been assumed that a magnetic disk archive can accommodate this level of archive access without requiring any additional hardware or software.

The following table (Figure 4) calculates the average access for LTO3 and UDO2 drives by accounting for all the steps in a full media exchange procedure. 2,500 requests per day equates to 313 requests per hour. The number of drives required to meet the access cycle for read operations is listed in the last row of the table. The LTO configuration requires a higher drive-media ratio to meet the same random access specification since load/unload and seek times are much slower than optical.

[2] Refer to section 5.1 for a full explanation of the drive count calculation.

Drive / Library Specs.	***UDO2***	***LTO3***
Load Time	5 sec	19 sec
Unload Time	3 sec	19 sec
Average Seek Time	35-50 msec	46 sec
Average Rewind Time	0 sec	44 sec
Media Exchange Time	6 sec	6 sec
Average Data Access	5 sec	5 sec
Average Access Cycle	19 sec	139 sec
Access Cycles per Hour	189 cycles	26 cycles
Drive Count for Read	2 drives	12 drives

Figure 4 – Drive Access Cycles for Reading

All performance data listed in Figure 4 was taken directly from the drive and library specifications of each vendor. For the sake of the analysis, worst case random access is assumed which means that each access request will require loading a new piece of media. These calculations also ignore the integrated disk cache that part of the architecture of the UDO Archive Appliance. The Archive Appliance cache buffers both read and write requests to the archive, reducing demand on the UDO drives.

If access patterns are more predictable or if physically adjacent files are often accessed together, the need for media exchange could be reduced. Also in this scenario the files being accessed are small, requiring very little time to read the files once located. In environments where files are larger, the drives with the higher streaming data rates (i.e. tape) will realize improved performance.

In the case of the LTO configuration, it may not be practical to provide a 40TB library with 12 drives, but since the LTO media capacity is so much larger than optical it can be assumed that tape exchanges may be less frequent. For the sake of this report, the number of LTO drives required for read operations has been set at 10 despite the calculations above. In practice, this may or may not prove to be sufficient drive resource depending on the actual library access pattern.

In addition to the drives required to read data from the archive, one drive has been added to each automated library to support archive write requirements. A single LTO3 or UDO2 drive is sufficient to write the required 8GB to 10GB of new data archived each day. One additional drive has also been added to each configuration for overall system redundancy in the event of drive failure. Figure 5 summarizes the total required drive count.

Drive Type	***Read Drive Count***	***Write Drive Count***	***Spare Drive Count***	***Total Drive Count***
LTO3	8	1	1	10
UDO2	2	1	1	4

Figure 5 – Total Drive Count

5.2 Software Configurations

In order to fairly compare the four different archive systems, it is necessary that each solution to have roughly the same interface and management environment. Plasmon's UDO Archive Appliance and the NetApp's FAS 3020 have a standard, network attached, filesystem interface that provides access to the archive. In the case of the FAS 3020, the NetApp SnapLock software was added to the configuration to provide WORM emulation to address the customer need for record authenticity. No additional compliance software is required for the Archive Appliance because it employs true WORM media for archiving data.

To approximate this same interface with the EMC Centera, Universal Access hardware and software were added to the configuration to provide a filesystem interface and their Government Edition software is required for WORM compliance.

Additional third party software is required for the Quantum i500 so it can be presented as a filesystem. In this example, QStar HSM was selected. QStar HSM is a well-respected archival management and HSM application for enterprise environments and will fully meet the archival requirements of this scenario. A Windows based server has also been added to the i500 configuration since the SCSI tape library requires a server for connectivity and for the installation of the QStar software.

Refer to Section 10 for more details on the Centera software configuration selected for this report.

6 TCO Component Analysis

The TCO analysis is based on quantifiable expenses for a 40TB data archive over the first three years of operation. Costs that could not be fully quantified, are not statistically significant or are subject to interpretation were not included. The cost components included in the analysis are:

- Hardware Acquisition
- Software Acquisition
- Initial Media Acquisition
- Hardware Maintenance
- Software Maintenance
- Floor Space
- Power and Cooling

A detailed spreadsheet of all figures and calculations can be found in Section 9. Section 10 provides additional details on the Centera hardware and software configuration. All figures are list prices in US$ taken directly from vendors' price lists. No special pricing or discounts have been used.

The financial model assumes full archive capacity from the first day of operation. While it is true that all of these solutions can all be implemented with smaller configurations and expanded over time, each scales in different ways. Consequently, it would have been too complex to fairly represent the differences in scalability and was not essential for the purposes of this analysis.

6.1 Capacity Adjustment Calculation

As shown in Figure 3, it was not practical to match all systems to a 40TB archive capacity. In order to provide an exact capacity match, it would have been necessary to select larger configurations than required and partially "depopulate" them to meet the 40TB target. This technique would be subject to manipulation and could put some systems at a significant financial disadvantage.

This analysis has, therefore, chosen to select fully populated systems as close to 40TB as possible, calculate the $/GB of each and adjust the overall system cost to match the 40TB target capacity. While this means that the adjusted system cost does not reflect an exact configuration, it provides a fair method of comparing relative costs of the different technologies. Figure 6 summarizes the adjusted cost using the technique described above.

Archive Type	Actual System Capacity	Adjusted Cost for 40TB (US$)	Adjusted Cost per GB (US$)
i500	51.2TB	209,049	5.23
AA638	38.3TB	211,337	5.28
FAS 3020	43.5TB	446,266	11.10
Centera Parity	42.5TB	509,709	12.74
Centera Mirror	41.4TB	809,518	20.24

Figure 6 – Capacity Adjustment Calculation

6.2 Results Overview

Figure 7 summarizes the findings of the TCO analysis in a graphical format using adjusted system costs over three years of operation. Since the final numbers do not reflect all possible costs, the absolute values are not complete. However, the objective of the analysis is to offer an even-handed assessment of all major expenses and to provide an accurate relative cost for each of the selected configurations.

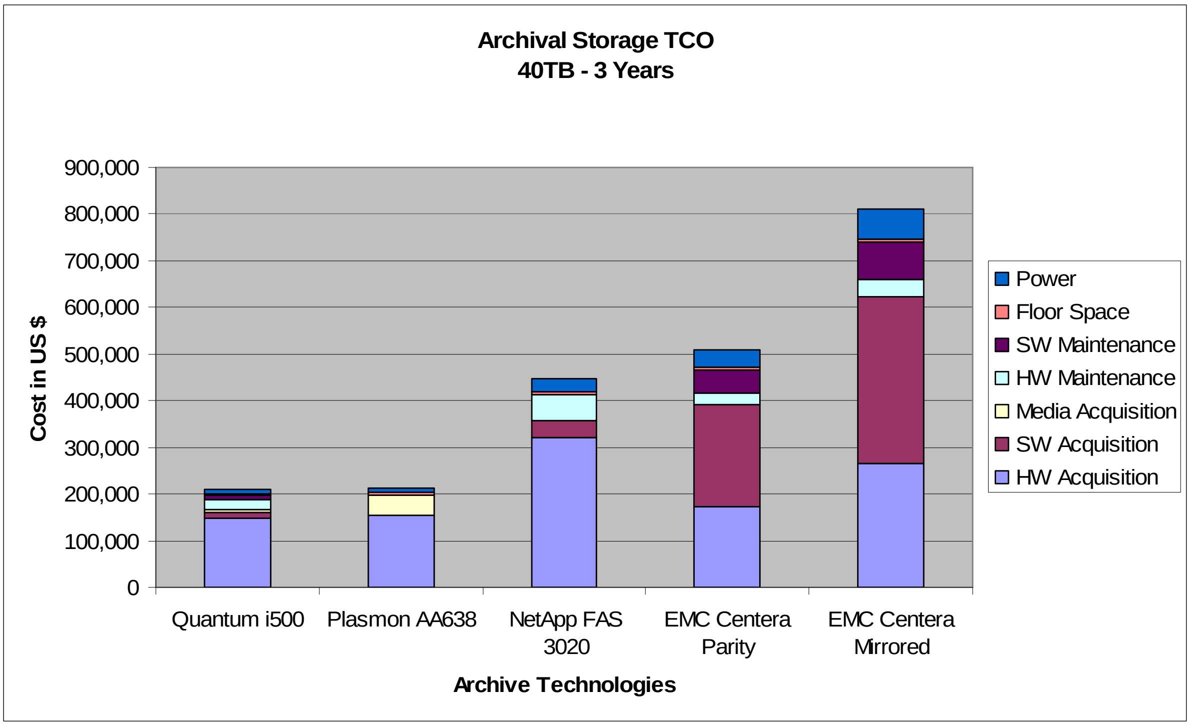

Figure 7 – Archival Storage TCO Detailed Summary

The results demonstrate that TCO for a 40TB archive is lowest with the Quantum i500 LTO3 library and the Plasmon's UDO Archive Appliance. The cost of these two configurations is virtually identical.

While the TCO for the NetApp FAS 3020 and the Centera Parity configurations are very close, they are more than twice the price of the Quantum and Plasmon solutions. The Centera Mirrored configuration is by far the most expensive option. It is nearly 4x the cost of the Quantum and Plasmon configurations. Figure 8 provides a complete cost ratio analysis for each of the configurations so that relative cost for a given configuration can be more easily compared.

Archive Type	*i500*	*AA638*	*FAS 3020*	*Centera Parity*	*Centera Mirror*
i500	1.00	1.01	2.13	2.44	3.87
AA638	0.99	1.00	2.10	2.41	3.83
FAS 3020	0.47	0.48	1.00	1.15	1.82
Centera Parity	0.41	0.41	0.87	1.00	1.59
Centera Mirror	0.26	0.26	0.55	0.63	1.00

Figure 8 – Cost Ratio Analysis

Figure 9 is a summary of all the costs of the TCO components using the adjusted system cost over three years. Figure 10 provides a breakdown on the relative percentage of the individual TCO components also using adjusted system costs. The two summaries are useful in analyzing the distribution of costs across all configurations and are referenced in the subsequent sections.

It should be noted that it is becoming difficult to break out the individual component costs as vendors increasingly bundle hardware, software and maintenance into a single price. For example, there is no specific line item for software cost with the Plasmon AA638, and three years of warranty and maintenance are also included. This can skew the analysis since all these costs are included in the hardware price. This analysis breaks out individual costs where they are available.

Archive Type	HW$	SW$	Media$	HW Maint$	SW Maint$	Floor Space$	Power$	Total$
i500	146,906	12,969	7,559	21,368	7,003	3,386	9,858	209,049
AA638	154,047	0	41,978	0	0	8,708	6,482	211,337
FAS 3020	321,523	37,517	0	53,650	0	3,386	27,772	443,847
Centera Parity	171,576	219,859	0	23,164	50,160	5,805	39,145	509,709
Centera Mirror	264,734	359,034	0	36,035	79,614	5,805	64,269	809,518

Figure 9 – TCO Adjusted Component Costs

Archive Type	HW%	SW%	Media%	HW Maint%	SW Maint%	Floor Space%	Power%
i500	70.3	6.2	3.6	10.2	3.3	1.6	4.7
AA638	72.9	0.0	19.9	0.0	0.0	4.1	3.1
FAS 3020	72.4	8.5	0.0	12.1	0.0	0.8	6.3
Centera Parity	33.7	43.1	0.0	4.5	9.8	1.1	7.7
Centera Mirror	32.7	44.4	0.0	4.5	9.8	0.7	7.9

Figure 10 – TCO Adjusted Component Cost Percentages

6.3 Hardware and Software Acquisition

Hardware acquisition costs for i500, AA638 and FAS 3020 are the highest, averaging 72%. The hardware percentage for Centera is lower at 33%, but is more expensive than the i500 and AA638 in real terms.

Regarding software acquisition, the software component for the i500 and FAS 3020 average about 7% of the total system cost. Software is included in the cost of the AA638 solution so cannot be broken out separately. The most surprising software percentages are with Centera. Software accounts for approximately 44% (on average) of the total system cost and is more than 27 times the price of the least expensive software. Centera software is one of the most expensive components for any of the archive systems.

6.4 Media Acquisition

Since the Centera and FAS 3020 do not make use of removable media, this component does not factor into their overall percentages. This could partially explain why the hardware percentage on Centera is lower than that of the library configurations.

Due to the higher 400GB capacity of LTO3 tape, it offers the least expensive removable media cost at $7,559. The lower capacity of 60GB UDO2 requires more pieces of media resulting in a much higher media cost of $41,978.

6.5 Hardware and Software Maintenance

The Centera system has the highest actual cost and similar overall percentage for hardware and software maintenance than the other configurations. The model demonstrates that annual maintenance contracts for the first three years on a Centera system can be as high as $38,000, much higher that that of the other technologies. These numbers are also misleading because the first two years of hardware maintenance are included in the cost of a Centera system. This means that after year two the on-going cost of hardware maintenance will jump sharply. This is an important consideration that should not be overlooked for any archival storage system designed to be operated for many years.

Hardware and software maintenance cost is much lower for the other archive systems. It is included in the cost of the AA638 so cannot be easily assessed until it is renewed after the first three years.

6.6 Floor Space

In this model a cost of $3,235 per m^2 per year was used. While the cost of office space can vary dramatically from city to city, this estimate is representative for computer room space in a large North American city.

The cost of floor space turns out to be a relatively small and fairly equal expense for all the configurations. Overall, floor space is a modest cost that does not contribute significantly to the TCO analysis. It should be noted that the percentage floor space cost for the AA638 appears much higher than the other systems. While the AA638 is physically larger, the percentage of system cost is disproportionately higher because the AA638 has fewer components that make up the total system cost. As a result, each component takes on a higher weighting.

6.7 Power and Cooling

The cost of commercial power is set at $0.2/KWH in this model with an 8% year-on-year increase in energy cost. This figure provides a reasonable average for US, European and Asian power rates but in some cases may not reflect regional increases in electricity cost. The model calculates an additional NCPI (Network Critical Physical Infrastructure) power consumption requirement using a 1.25 System Power ratio. NCPI power includes: air conditioning, circulation and humidification, transformation and UPS overheads. All power numbers were drawn directly from published vendor specifications.. Figure 11 provides a summary for the figures used in calculating power and cooling cost over three years of operation.

Archive Type	System Power (W)	NCPI Power (W)	Total Power (W)	Adjusted Total Power (W)	Annual Power (KW)	3 Year Cost ($)
i500	986	1,233	2,219	1,733	3,037	9,858
AA638	485	606	1,091	1,140	1,997	6,482
FAS 3020	2,360	2,950	5,310	4,883	8,555	27,772
Centera Parity	3,250	4.063	7,313	6,882	12,058	39,145
Centera Mirrored	5,200	6,500	11,700	11,304	19,805	64,296

Figure 11 – Power Consumption Costs

Since both the Quantum i500 and the Plasmon AA638 use automated library technology rather than spinning magnetic disk, their energy costs are dramatically less expensive to operate than NetApp or EMC solutions. In the worst case, even the Low Power model of the EMC Centera system can cost over $20,000 each year for power consumption alone.

Energy consumption is becoming a much more important consideration for CIOs when making key acquisition decisions. Many companies are now looking for ways to reduce power consumption by eliminating power hungry equipment. The use of low power consumption library architectures for archive data has very distinct financial and environmental advantages.

7 Additional Considerations

This report has included only quantifiable costs in the calculation of TCO for the different archive solutions. However, there are a number of less quantifiable considerations that the authors of this report feel are important to note. These issues can have a significant impact on TCO calculations so may need to be factored in when building TCO models for specific corporate applications.

7.1 System Administration

The TCO model has not included the expense of system administration. Estimating the cost of data archive management can be subjective and is dependant on the skill sets of the IT organization. In addition, it can be reasonably demonstrated that the administration overhead for all the technology types considered is approximately the same (short-term and long-term). Some vendors have suggested that removable media configurations require dramatically higher administration resource than magnetic disk based systems. This is simply not a valid assumption, as it cannot be substantiated by end user experience. For these reasons and since this report is tasked with identifying cost differences between the competing technologies, administration costs were not included.

7.2 Data Migration and Hardware Upgrades

Data archives must be designed to operate for years/decades, which means that data with long-term retention requirements will need to be migrated periodically from older to newer technologies. This TCO analysis accounts only for the first three years of operation and it can be assumed that no migration or hardware upgrades will be necessary during this period. Since tape, disk and UDO technologies all have different media and system life spans, it is important to understand how frequently data migration and upgrades are required in order to ensure long-term data availability and to assess ongoing operating costs. Stable, long-life UDO media offers a distinct advantage over tape and disk in terms of reducing migration frequency.

7.3 Tape Media Maintenance

When storing data on magnetic tape for extended periods of time media manufacturers and industry bodies often recommended that tape media be subjected to a maintenance program in order to ensure long-term data availability. Though the exact requirements for a tape maintenance program are up for debate, the requirement cannot be ignored when considering tape for a long-term archive.

Tapes that have not been accessed for an extended period should be “re-tensioned” and tape error rates monitored to identify ageing tapes that may require refreshing to new media. It is up to the individual organization to determine the value of their data and the frequency of management when using tape for long-term archival storage. There are both media and considerable administration costs associated with tape maintenance, in addition to incremental library resources needed for re-tension and refresh operations. This type of proactive media maintenance is not required for UDO or redundant magnetic disk systems.

7.4 Data Authenticity and Audit Trails

Data authenticity can be a critical consideration for many archival storage environments and the use of WORM media is an important piece in the data authenticity “chain of trust”. Magnetic disk and tape technology are inherently rewritable, but offer WORM emulation implemented through firmware/software that prevents rewritable media from being overwritten. Only UDO phase change optical media provide true Write Once data recording. The use of true Write Once media can reduce the cost of audit trail management and should be considered “best case” for archive environments where data authenticity is a high priority.

7.5 Off-line Secondary Media Copies

Regardless of the stability of the media, it is important to retain more than one copy of all valuable information in order to protect against disasters or site failure. Often organizations look to store redundant data sets off-site in a protected vault to minimize cost. When using removable tape or UDO media, it is a reasonably simple and cost effective process to make a second copy of media for storage in a vault.

When using a magnetic disk archive such as the EMC or NetApp solutions, it is not possible to vault the magnetic disk. As a result, it is necessary to implement a tape or UDO strategy with extra cost software to create the second data copy for off-line vaulting. This adds a great deal of additional cost and complexity to the archive infrastructure. Alternatively, a second mirrored disk subsystem could be installed to perform remote replication. While this may meet requirements for protection against site failure, it is an extremely expensive strategy that can be very difficult to justify in environments where archive data is accessed infrequently.

7.6 Installation Considerations

The installation of storage hardware and application software is a routine procedure for many IT managers. Proper planning and training will ensure a quick and successful installation process. In the case of Centera and other magnetic disk archive systems, it is important to mention two considerations that are somewhat out of the ordinary: the weight of the hardware and the power connection. As an example, a fully populated Centera cabinet weighs nearly 700kg (1500 lbs). Structural reinforcement may be necessary to ensure that the weight can be safely supported. In addition, the Centera power lead requires a heavy-duty electrical connection (L6-30R / IEC-309-332R6) for the demanding power requirements of the cabinet. This may require special electrical infrastructure before a disk based archive system can be installed. Weight and power availability is also a consideration, for the NetApp FAS 3020 configuration.

8 TCO Analysis Summary

The figures provided in this Archival Storage TCO Analysis do not attempt to take into account all possible costs of a professional archive environment since many expenses are application or site specific. However, the analysis does address the most significant costs and seeks to provide a valid and representative ranking for each selected configuration.

The primary cost differences have less to do with the storage technology and more to do with their specific implementations. Indeed, the Quantum LTO3 tape configuration and the Centera products are both based on a magnetic storage, but stand at extreme ends of the TCO spectrum. Overall, the Plasmon UDO Archive Appliance and Quantum i500 offer a dramatically lower TCO proposition than NetApp or Centera alternatives and require much lower on-going operating expenses.

This report provides quantifiable insight into archive acquisition and maintenance costs, but it is also important to balance the financial analysis against other considerations. Business and technical issues such as regulatory compliance, data authenticity, media longevity and performance are equally valuable metrics and must be considered when evaluating archival storage strategies. Readers are encouraged to use the results of this report as a starting point for their own archival storage TCO evaluations.

9 TCO Summary Spreadsheet

			LTO Scalar i500 10 LTO3 Drs	UDO AA638 4 UDO Drs	Disk NetApp FAS 3020	Disk Centera Parity	Disk Centera Mirror
Capacity in TB			**51.2**	**38.3**	**43.5**	**42.5**	**41.4**
Hardware Cost							
Storage Hardware			183,039	147,500	349,656	182,300	274,000
Server Hardware			5,000				
Total			188,039	147,500	349,656	182,300	274,000
% of Total Cost			*75.03*	*78.59*	*77.91*	*36.92*	*35.80*
Software Cost							
QStar HSM			16,600				
NetApp SnapLock					40,800		
EMC Centera SW						233,600	371,600
Total			16,600	0	40,800	233,600	371,600
% of Total Cost			6.62	0.00	9.09	47.31	48.56
Media Cost	**Unit Cost**	**Slot Count**					
400GB LTO3 WORM	118	82	9,676				
60GB UDO2	63	638		40,194			
Total			9,676	40,194	0	0	0
% of Total Cost			*3.86*	*21.41*	*0.00*	*0.00*	*0.00*
Hardware Maintenance	**3 Years**		27,351	0	58,344	24,612	37,296
% of Total Cost			*10.91*	*0.00*	*13.00*	*4.98*	*4.87*
Software Maintenance	**3 Years**		8,964			53,295	82,401
% of Total Cost			*3.58*	*0.00*	*0.00*	*10.79*	*10.77*
Floor Space Cost	**m^2**	**Yearly $/m²**					
Scalar i500 - LTO3	0.35	3,225	3,386				
AA638	0.90	3,225		8,708			
NetApp	0.35	3,225			3,386		
Centera	0.60	3,225				5,805	5,805
Total			3,386	8,708	3,386	5,805	5,805
% of Total Cost			*1.35*	*4.64*	*0.75*	*1.18*	*0.76*
Power and Cooling			9,858	6,482	27,772	39,145	64,296
% of Total Cost			*3.93*	*3.45*	*6.19*	*7.93*	*8.40*
Actual System Cost			250,630	187,694	448,800	493,807	765,297
Cost per GB			4.90	4.90	10.32	11.62	18.49
Adjusted to 40TB			195,805	196,025	412,690	464,760	739,417
Total Adjusted Cost			209,049	211,215	443,848	509,709	809,518
Adjusted Cost per GB			5.23	5.28	11.10	12.74	20.24

Figure 12 – TCO Summary Spreadsheet

10 Centera Pricing Details

Centera Hardware and Software Pricing

US$ List Prices

20 Node Partity Configuration - 42.5TB

Line	Qty	Product	Description	Unit List	Total List	Comment
1	1	CNRRK	40U T RACK SP POWER	4,200	4,200	19" Rack
2	1	CAT6CBL100	UTP CAT6 CABLE 100 WHITE	200	200	Rack Component
3	1	CNR-CBL15-G4	SHORT UPLINK CBL KIT	100	100	Rack Component
4	1	CNR1PH	SP PWR EXP EMC RK	1,200	1,200	Rack Component
5	1	CNR4N5MBAG4	4 NODE 750GB MANF BASE G4	48,800	48,800	4 Node Base Unit
6	4	CNR4N5MEXG4	4 NODE 750GB MANF EXPAN G4	30,200	120,800	4 Node Expansion Unit
7	1	CNRCUA1NG4	SINGLE CUA NODE G4	6,200	6,200	Centera Universal Access Node
8	2	CNRMODEMG4	CNR MODEM GEN 4	0	0	Modem
9	1	PW40U-IEC3	DUAL 40U PWRCORD-INTL	800	800	Dual Powercord
					182,300	**Total Hardware Cost**
10	1	CNRGPLLIC	CNR GEN PUBLIC LIC SW	0	0	Centera General Software License
11	1	CNRCONSOLESW	CONSOLE SW LICENSE	0	0	Centera Console License Software
12	5	CNR4NGOVEDLIC	4NODE GOV EDITION SW LIC	6,200	31,000	Government Compliance Software (1 license / 4 node set)
13	5	CNR4NPSWG4	PAR SW RTU LIC	38,200	191,000	Parity Software (1 license / 4 node set)
14	1	CNRCUAESW	CNR EXT DELL CUA SW	11,600	11,600	Access Software
					233,600	**Total Software Cost**
15	1	WU-PREHW-001	PREMIUM HARDWARE SUPPORT	24,612	24,612	3 year premium HW service (20 nodes)
					24,612	**Total Hardware Maintenance**
16	1	M-PRESW-001	PREMIUM SOFTWARE SUPPORT	53,295	53,295	3 years combined software maintenance (20 nodes)
					53,295	**Total Software Maintenance**
					493,807	**Total System Cost**

32 Node Mirror Configuration - 41.4TB

Line	Qty	Product	Description	Unit List	Total List	Comment
1	1	CNRRK	40U T RACK SP POWER	4,200	4,200	19" Rack
2	2	CAT6CBL100	UTP CAT6 CABLE 100 WHITE	200	400	Rack Component
4	2	CNR-CBL15-G4	SHORT UPLINK CBL KIT	100	200	Rack Component
5	1	CNR1PH	SP PWR EXP EMC RK	1,200	1,200	Rack Component
9	1	CNR4N5MBAG4	4 NODE 750GB MANF BASE G4	48,800	48,800	4 Node Base Unit
10	7	CNR4N5MEXG4	4NODE 750GB MANF EXPAN G4	30,200	211,400	4 Node Expansion Unit
11	1	CNRCUA1NG4	SINGLE CUA NODE G4	6,200	6,200	Centera Universal Access Node
12	2	CNRMODEMG4	CNR MODEM GEN 4	0	0	Modem
13	2	PW40U-IEC3	DUAL 40U PWRCORD-INTL	800	1,600	Dual Powercord
					274,000	**Total Hardware Cost**
14	1	CNRGPLLIC	CNR GEN PUBLIC LIC SW	0	0	Centera General Software License
15	1	CNRCONSOLESW	CONSOLE SW LICENSE	0	0	Centera Console License Software
18	8	CNR4NGOVEDLIC	4NODE GOV EDITION SW LIC	6,200	49,600	Government Compliance Software (1 license / 4 node set)
16	8	CNR4NMSWG4	MIRR SW RTU LIC	38,800	310,400	Mirror Software (1 License / 4 node set)
19	1	CNRCUAESW	CNR EXT DELL CUA SW	11,600	11,600	Access Software
					371,600	**Total Software Cost**
20	1	M-PREHW-001	PREMIUM HARDWARE SUPPORT	37,296	37,296	3 year premium HW service (32 nodes)
					37,296	**Total Hardware Maintenance Cost**
22	1	M-PRESW-001	PREMIUM SOFTWARE SUPPORT	82,401	82,401	3 years combined software maintenance (32 nodes)
					82,401	**Total Software Maintenance**
					765,297	**Total System Cost**

Figure 13 – Centera Pricing Details

Plasmon

370 Interlocken Boulevard, Suite 600, Broomfield, CO 80021 U.S.A. Tel: 0800.451.6845

Whiting Way, Melbourn, Nr. Royston, Hertfordshire SG8 6EN, U.K. Tel: +44 (0)1763 262963

www.plasmon.com