Table of Contents

Overview of the Placement

My first work term lasted four months and took place between May, 1998 and August, 1998. I was fortunate enough to be hired by IBM Canada Ltd. in Markham, where I held the position of Storage Systems Support. I was hired into the Storage Systems Division of the large IBM corporate structure, and within this division I worked in the Canadian Sales and Marketing department. My manager was Graham Whillier, who is the National Sales Manager for Storage Systems in Canada.

Since IBM's beginning in 1924, IBM has become synonymous with computers. IBM has been at the forefront of computer technology for many years and today their role in the computer industry is targeted towards developing and manufacturing computer systems, software, networking systems, storage devices and microelectronics. These are in addition to providing consulting services, systems integration, and technical support. As the computer industry has changed over the years so has IBM. Today IBM is known to the computer world as:

1. Creating the industry's most advanced information technologies
2. Helping customers apply those technologies to improve what they do and how they do it *(http://www.ibm.com/IBM/)*

IBM Storage Systems Division has been one of IBM's most successful divisions. The Storage Systems Division (SSD) has been leading and revolutionizing the storage industry since its first hard disk drive in 1956. SSD continues its tradition in the manufacturing and distribution of hard disk drives, storage subsystems, tape and optical media, storage management software and other essential storage components.

SSD has grown internationally over the past years, including penetration into the Canadian market. The Canadian Sales and Marketing department is involved in the marketing, support and integration of its products here in Canada.

There are numerous sales representatives and Information Technology specialists across Canada with the main office located in Markham. At this location the department is divided into employees focusing on one of three system types: High-End (mainframes), Mid-Range (AS/400) or Open Systems

(PC systems such as UNIX and Windows NT). Within each system type there is another breakdown of employees geared towards each storage product, such as tape or disk, and they are responsible for aspects relating to that specific product's usability and connectivity towards the system type they know most about.

The sales representatives are in charge mainly of the presentation and marketing of the products. The Information Technology specialists are specialized in specific storage products and are responsible for the marketing of the product they specialize in. In addition, they are also responsible for any pre-sales and (some) post-sales support, in particular technical questions and the integration of their product with a specific system type.

My main role in the department was in the position of an Information Technology specialist, even though my title was Storage Systems Support. A fellow internship student from the University of Western Ontario ("UWO") and I were the Canadian contacts for the client/server storage management software called *ADSTAR Distributed Storage Manager (ADSM)* which encompasses all three system types mentioned above, however I worked mainly in Open Systems.

My responsibilities included answering customer and employee questions about the product, which often required research into the product and usually related to the usability and connectivity of the product on one of the many system platforms it covers. I was also responsible for setting up demonstrations of products to be shown at customer meetings and the preparation of any sales material needed for these meetings. I also setup a test ADSM server on Windows NT for the department to use, and later this server was the testing ground for my colleague and me to learn more about the product.

In addition to my work with ADSM, I also helped out the department with any technical problems relating to their systems, such as printer or network problems. I watched “over the shoulder” of some of the Information Technology specialists working with IBM's Serial Storage Architecture (SSA) adapters, which involved disk subsystems. My work with SSA also included updating microcode on the adapters through Windows NT and IBM AIX, which were used in the department's demonstration program. I was fortunate enough to be allowed by my manager to go to several IBM Training and Education courses on ADSM, AIX and Lotus software. This allowed me to learn more about the UNIX operating system which I believe, will be a great benefit to me in the future.

Introduction

During my work term with IBM, I spent the majority of my time working with the client/server storage management software called *ADSTAR Distributed Storage Manager* or *ADSM* for short. Initially I was assigned to a user-group that talked about the new features of the version 3 server which was released in February, 1998. Afterwards I spent time talking with the Information Specialist on ADSM. Her name was Betty Cho and she had been working with IBM for almost thirty years. Her knowledge about computers and networks, as well as ADSM was overwhelming. Within one day she had taught me the basic principles behind the product. Unfortunately, the time I spent with Betty was limited due to the fact that she was about to transfer to a new position in the company. This left an open position within the department and I, in conjunction with my fellow intern student from UWO, were the department's temporary replacement solution. This was a daunting task for us, since both of us had limited exposure to the computer industry let alone an award winning software package.

From the second day onward we spent two months self-training on ADSM. Within two weeks I had read approximately fifteen hundred pages of reference and technical manuals relating to ADSM. We both were sent on a week long course that talked about solution design and implementation. We also setup a demonstration ADSM server on Windows NT and a test ADSM server for our department to use, also on Windows NT.

By the third month we had gained enough knowledge about the product that we were open for demonstrations and questions about the product. So as one can see, my work term was mainly focused on ADSM. It would then naturally make sense for me to share some of the information I learnt about the product, and hopefully I can give the reader a sense of why ADSM is recognized in the industry as an award winning storage management solution.

ADSTAR Distributed Storage Manager

What is ADSM?

As mentioned earlier ADSTAR Distributed Storage Manager (ADSM) is a client/server storage management software package/solution. To properly describe and understand ADSM it is appropriate to first review a brief description of the storage management industry.

A *storage management solution* "should provide functions in all key areas of storage management, data protection, data retention, disaster recovery and hierarchical storage management" *(http://www.storage.ibm.com/software/adsm/adwhcoo.htm)*. Almost every computer owner has gone through the difficulty of losing a file and wishing a back up existed. This scenario sets the scene for storage management. Recently the Information Technology (IT) industry has grown at an amazing rate and corporate data is doubling in most cases. With most companies now moving towards larger and more complex computer systems due to decreased network and hardware costs, increased focus has been given to centralizing storage. With these new trends arising, loss of data (see *Appendix A*) could be devastating to a company. According to PC Week, "Fifty percent of the companies that lose critical business systems for 10 or more days never recover. Ever." *(PC Week, February 6, 1995, page 129)* Data critical to the operation of the company must therefore be stored in a safe and manageable way so that access to the data can be quick and reliable.

With ADSM, all the points mentioned in what a storage management solution traditionally needed to be have been addressed. As well, ADSM has been designed to cope with these new trends that are developing in the IT industry. A typical ADSM configuration can be seen on the next page. It consists of a central ADSM server which has some sort of storage hierarchy attached to it such as disk, tape or optical (or any combination of them). Attached to the ADSM server through one of nine different communication protocols are ADSM clients, which can be personal computers and/or LAN file servers, and databases and applications.

Figure 1. A typical ADSM configuration. *(http://www.storage.ibm.com/software/adsm/adwhchoo.htm)*

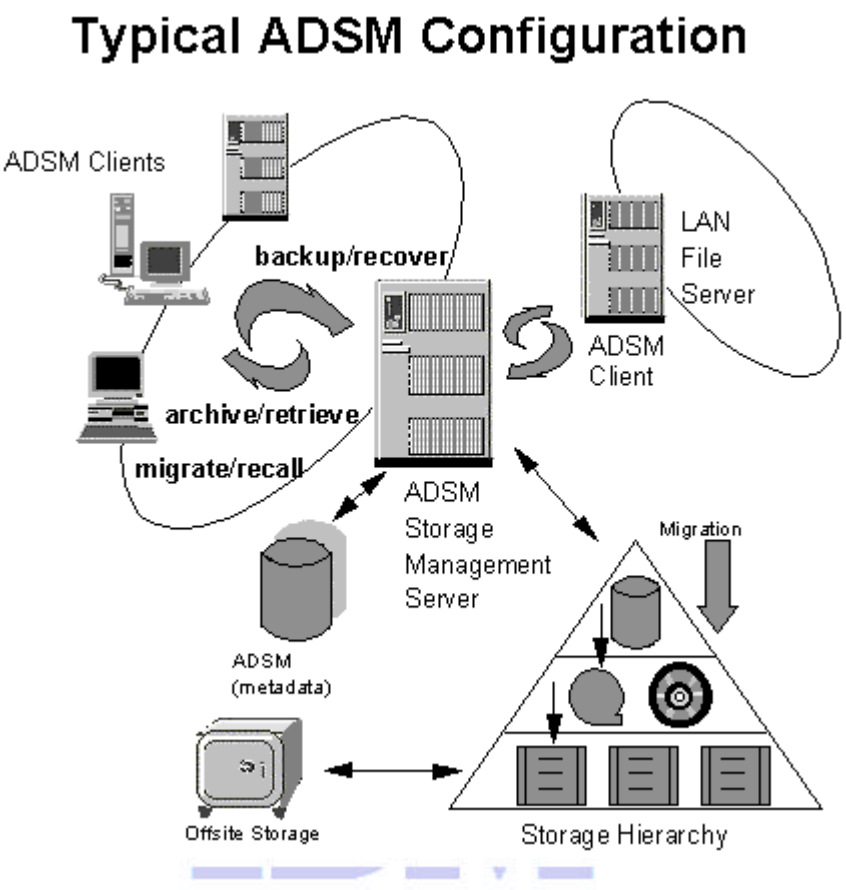

The ADSM server is available for nine different platforms and the ADSM client is compliant with over thirty different platforms (see *Appendix B*). ADSM also allows certain databases and applications to interact with the functions of the ADSM server. Sometimes this is done through the database or application if it contains its own backup and restore functions and sometimes this is done through what is called an ADSM Connect Agent (see *Appendix B* for listing of Connect Agents). A Connect Agent allows communication with the database or application with the ADSM server through the ADSM application-programming interface (API).

Main Components

The ADSM backup/archive client runs on the workstation and, depending on the platform, provides both a graphical user interface (GUI) and command line interface (CLI), and allows a user to back-up, restore, archive and retrieve data files. A feature of the client is file compression, which can reduce network traffic and the amount of storage on the server when backing up or archiving a file. Through the client ADSM allows for cross-user restore and cross-platform restore. Cross-user restore allows one user to allow someone else to restore their files, while cross-platform restore allows a user to

restore files from one platform to a different platform, provided the platforms share the same file structures, such as OS/2 to DOS, but not AIX to DOS.

Figure 2. ADSM client retrieve screen shot (*http://www.storage.ibm.com/software/adsm/adtc2p1.htm)*

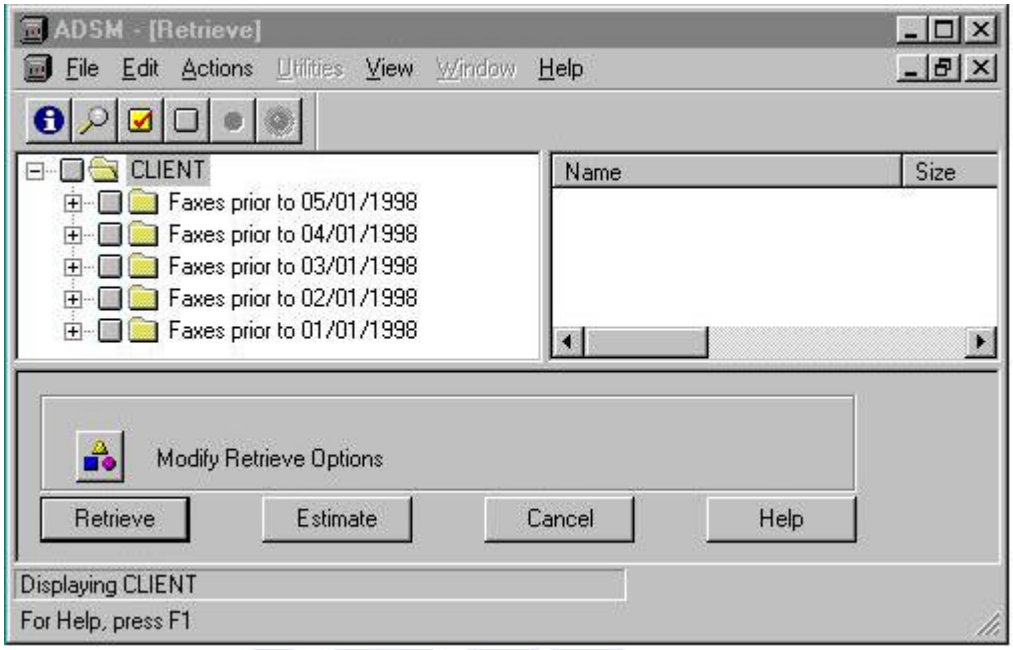

The ADSM administrative client allows an ADSM administrator to control and monitor the server activity, define storage management policies for workstation files and setup schedules to provide backup and archive services. The administrative client also comes in both a GUI and CLI.

The ADSM server provides storage resources and services for the clients. Users can backup or archive their files onto server storage resources such as disk, tape or optical devices, which could be defined in a hierarchy of storage, that the ADSM server manages and monitors. There are two key components of the ADSM server. One is the storage pools, which are where the client files are stored and second is the database. The database serves as an inventory or index to the client files within the storage pools, it also holds all the information about the client nodes, schedules and an activity log of all the actions relating to the server, just to name a few. Within the database there is a recovery log that keeps track of all changes made to the database. If there is a hierarchy of storage present, data can be migrated/moved down the hierarchy to less expensive media. Additional management functions are provided, such as reclamation (re-using of tapes) and collocation (bringing together similar data from different tapes into one area) for tape management so that a company's media supply is being used

efficiently. The ADSM server allows for concurrent backups so there is no limit to the number of users at one time.

Analysis of ADSM

In this section I will discuss some of ADSM's major functions and analyze the effectiveness, efficiency and usefulness of their tasks in respect to the needs of the storage management industry.

Backup/Restore and Archive/Retrieve

ADSM's main function is to be able to backup and restore a client's file systems. The backup process creates a copy of a client file on the ADSM server. ADSM allows for three different types of backups: full, incremental and selective. A full backup is when the entire file system is backed up at one time. An incremental backup sends to the server only the files that have changed since the last backup. The first incremental backup acts as a full backup. ADSM determines if a file has been changed by looking at one of the following: file size, date/time stamp, file owner, file group, file permissions or attribute change time. A selective backup allows the user to decide what files should be backed up. The files are backed up according to policies set out by the administrator. The policies determine how many backup versions should be kept in the ADSM storage pools, and how long to keep those versions and whether to back up files that are in use.

Restore is the process of copying a backup version from the server to the client. The system performs the restore for the user, such that the user does not have to call the administrator to request the restoration of the file. It can all be done from the backup/archive client.

ADSM's backup/restore function offers the users the capability of doing a true incremental backup. Earlier software packages always followed the full backup paradigm, which usually meant long backup windows and high network traffic. The paradigm was that every night a full backup of the file systems would be done; obviously hundreds of gigabyte file systems, which most companies have today, would take an enormous amount of time to backup. With ADSM's incremental backup, a new paradigm has been created which makes it easier. Night one would still require a full backup but the remaining weeknights would only require an incremental backup, since only those files that have changed from the night before would be considered. This tremendously cuts down a company's backup window and distinctly decreases network traffic.

Another key issue that improves performance is the way ADSM handles small and large files. In previous software packages, the time required to backup small files was tremendously slower compared to the time required to backup large files. It was determined that the overhead required to do a backup of a file either small or large was the same. The same information such as location and file information had to be recorded in the server database and the size of a file was irrelevant. So what ADSM does is what it calls "*small file aggregation*" and it reduces the performance gap between small and large files. In an example, we have five small files, files 1 through 5, that are assembled as one object on the ADSM server, while file 6, which is a large file, remains a single file. By grouping the files this way, there is a reduction in the number of data pointers and ADSM backup entries, which requires less overhead on reclamation, migration and storage pool backup operations, and with fewer objects to manage, data transfer is faster. This new storage process allows large objects to move through ADSM's large buffers quicker and cuts down on CPU utilization because less overhead is required in handling individual file entries (see *Appendix C*).

The final aspect of ADSM's backup/restore usefulness is its *fault tolerance* features. In many situations files become corrupt or the network crashes. It could be costly if a user were in the middle of a backup or restore when this happened. In most cases the user would have begin again, thus wasting his or her time. However, the new ADSM features solve this problem. In terms of backup, if ADSM is in the middle of a backup and it comes across a file that is unreadable, it has the knowledge to skip the file, log it as an error and continue with the rest of the backup. This saves the client time by avoiding the consequences of repeating the operation. In terms of restore, if a restore is being done and the network crashes, once the network is up again, ADSM realizes where it left off and continues with the restore, saving time and unnecessary network traffic.

The archive/retrieve function is used for the purposes of long term storage. The archive process creates a copy of a client file on the ADSM server. As with backup, archived files are managed using policies, but the concept of versioning is not used since multiple versions of the same file can be kept through different archive copies. The importance of archiving can be seen for companies that require records to be kept for long periods of time, such as for tax purposes. The main difference between backing up and archiving a file is that once a file is archived the original file can be deleted since the

archive version is expected to be retained for a long time. The retrieval of a file is the same as a restoration except it occurs on archived files.

Central Scheduling

ADSM central scheduling allows for the automation of client backup, archive, restore and retrieve, as well as ADSM server administrative operations. The central scheduler consists of client and server processes that work together to execute the scheduled tasks. The administrator is responsible for defining and maintaining the schedules and determining the priority of client nodes so those clients with more important data are given preferential treatment. The scheduler allows for automation, which these days is an asset. Less employee time is required to overlook the everyday maintenance of the server and since everything is done from one central location it is easier to keep track of the functions going on. ADSM also combines well with system monitoring and event logging software such as Tivoli, HP Openview and IBM NetView. This software allows the server to run on its own with no human monitoring and if an error should occur it sends an alert to an operator who can then take care of the problem. Also, reports can be generated using SQL queries and, with ADSM's Open Database Connectivity (ODBC), popular spreadsheet applications can also be used to better control the ADSM environment.

Policy Management

Since in a company there may be different departments that require different storage requirements, ADSM enables an administrator to manage the backup and archive process, down to the file level, according to policies established for each department. Policy management makes ADSM a true system-managed storage implementation. Policy management follows a hierarchy of policy domain, policy set, management class and copy group. A *policy domain* is a group of clients that are working according to the same set of policy needs. Each policy domain can contain one or more *policy sets*, which can have one or more management classes. Only one policy set can be active at one time. This could be used for seasonal requirements for example. A *management class* contains a backup and/or archive copy group or no copy group. A *copy group* is where one specifies the file parameters such as where the file will be located, what if the file is in use and the number of versions just to name a few. A copy group exists for the backup and archive functions and each carries its own file unique parameters.

The idea of policies is well defined for large companies that require different storage needs for different files and/or departments. To make backups and archiving easier, something called an include/exclude list can be determined for each client. This list allows the selection of only those files included in the list to be backed up or archived while those files in the exclude list will not be selected; by default all files are included but this can be modified. As one can see this concept of policy management is quite complicated and takes some time before it can be completely understood. This point is one weakness the industry has discovered with ADSM. The idea of policy management is correct, but only the product has implemented it a way that makes it difficult to understand due to the many different layers in the hierarchy. The theory behind why each level is there makes sense, but implementing such a setup is not straightforward.

Disaster Recovery

Looking at *Appendix A*, it is likely that at some point something will happen to a company's storage structure. In the case of a severe loss of data, ADSM has come out with an add-on product call Disaster Recovery Manager (DRM) that automatically generates a server disaster recovery plan. It also manages offsite disaster recovery media and stores client machine recovery information. The main things that would need to be backed up in preparation of a disaster would be the ADSM database, recovery log, storage pools and copy pools (copies of storage pools). The report generated by DRM gives a layout to what should be done to prepare and recover from a disaster. These days most companies are realizing that they require a disaster recovery plan. Their data is too important to be lost, as many companies in Montreal realized when the ice storm hit there last winter. Setting up such a plan is tedious, but this add-on function makes the whole process easier to accomplish and something most large companies are realizing and adding to their storage management requirements. The one deterrent of DRM is the fact that it is an add-on feature that comes at a cost of about $10,000 CDN. But if one thinks about it, losing your company's complete data information could cost millions more. For many companies it is a tough decision when they purchase ADSM.

Hierarchical Storage Management (HSM)

HSM automatically allows the migration of infrequently-accessed data from online storage media such as workstations or local server hard drives, to less expensive media, such as tape or optical. This ensures more space on the faster disk media and is also more cost efficient. When this is done a stub or pointer is left on the client system pointing to the file in the storage hierarchy. When the file is

needed it gets recalled from the storage hierarchy for the user to use, the process of which is transparent to the user. This concept provides a true storage management function that can help in keeping a company's storage costs intact. Unfortunately HSM is only provided natively on AIX and Sun Solaris and via third party offerings on Novell Netware and Windows NT.

General

In general ADSM is robust and a very scaleable storage management solution. Once ADSM is initially setup, the addition of extra client nodes comes with no complex procedures, just the addition of the node to the server and the installation of the client code on the workstation. This allows for unlimited clients to be connected to an ADSM server. As well, the addition of extra storage devices is as easy as connecting the device to the server and configuring them into the storage pools.

The drawback of having such as sophisticated storage management solution is that the initial installation of ADSM is very difficult and time consuming. Some installations may take over a week, due to the heavy overhead of setting up storage pools and policies, which may take a fair amount of time depending on the company. The rule of thumb is that if a company has over ten servers, then ADSM is the storage management solution of choice. Otherwise another less sophisticated product may be a better option. Also noteworthy is that people in the industry, as well as IBM, realize that in terms of performance, ADSM's backup function is slower than that of most other competitors, but their restore function is the fastest in the industry. This has been hotly debated and IBM states that it is when you need to restore a file that time is most critical. For example, one has a presentation in fifteen minutes and has lost the presentation file, which was backed up. This poor person will be hoping for the fastest recovery of this file as possible. How long it took to backup the file is irrelevant. For some in the industry this does not suffice, but like I mentioned it is a topic always open for debate.

ADSM and its Competitors in the Market

ADSM has a strong influence in the marketplace. It is expected that ADSM's 1998 North American revenue will be $40 million (US). This constitutes a large portion of the market. What makes ADSM so strong in the market is its availability for all three system types (High-end, Mid-range and Open systems), and in particular AS/400 and mainframe which most competitors do not support. This gives ADSM a large advantage over its competitors.

Where competition is quite fierce is in the Open Systems area, where such products as Cheyenne ARCserve, Legato Networker, and Seagate Backup Exec are possible alternatives to ADSM. Overall the top three storage management products on the market are IBM ADSM, Legato Networker and possibly Cheyenne ARCserve. Like most products, there is always some sort of flaw. The main flaw with the competitors' products is their lack of scalability in comparison to ADSM. Although the competitors are good options for small companies with few servers and workstations, they do not suffice if a company were to grow in the future as trends in the industry suggest.

Many industry reports have been done on storage management software and most conclude with the same information I have given above. ADSM has won several awards for the introduction of their version 2 product in 1994 and its maintenance and upgrading of it to version 3 in 1998. The highest accolade comes from the recognized research firm of the Gartner Group, which ranked ADSM number one out of twenty-one storage management products in December 1995. Since then ADSM has improved substantially in many of the top areas of storage management and is still one of the top products on the market.

The Future of ADSM

The future of ADSM looks good. With IT and storage trends rapidly increasing, the need of a quality storage management solution is key to a companies functionality. ADSM improved greatly on its version 2 release with a version 3 release, for many of its servers which is Y2K compliant, in February 1998. It has also realized the increased focus on the Internet and has come up with a Web GUI for the administrative and backup/archive client. The web administrative client allows for an administrator from anywhere in the company with a web browser to make any necessary changes to the ADSM server. The web client allows a user to have access to the main functions of the backup/archive client from anywhere from within the company with a web browser in order to make any necessary actions. Also in September 1998 IBM announced the introduction of a version 3 server and backup/archive client for AS/400, which is in large demand, as well as some other newer platform clients such as Microsoft Windows 98 and Data General DG/UX. A version 3 server for IBM VM might also be available in the near future. IBM is ever seeking to make new enhancements to ADSM in order to give its customers the best storage management solution available.

Conclusions

In conclusion, ADSM is a storage management solution that fits into the needs of today's quickly changing storage management industry. It is available for use in all three system types and is scalable enough to meet the growing trends of corporate data storage. ADSM's advantages greatly outweigh its disadvantages and easily cater to the needs of most companies. It fits the definition of a storage management solution defined earlier. I have just given the reader a glimpse of what ADSM and the storage management industry is all about. Hopefully I have been able to bring enough insight into the situations faced by the storage management industry and have shown why ADSM is recognized as an award winning storage management solution.

Appendix A: Reasons for Data Loss

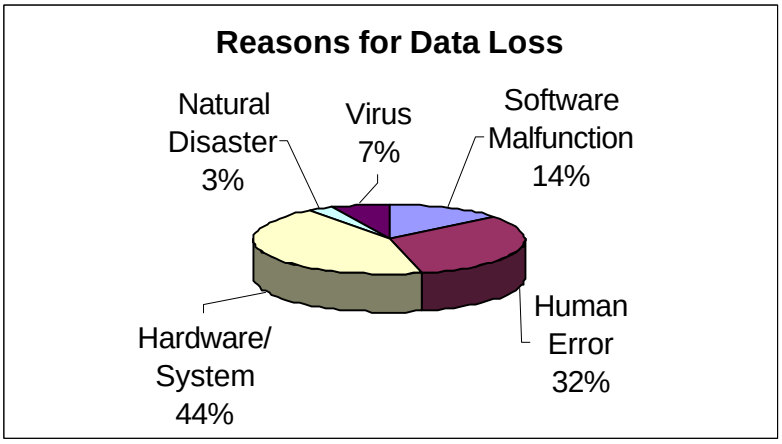

Figure 3. Reasons for Data Loss *(IBM Storage Roadshow - Toronto 1998)*

Appendix B: Server & Client Platforms and Connect Agents

Server Platforms

IBM OS/2
Microsoft Windows NT Intel
IBM AIX
Hewlett-Packard HP-UX
Sun Microsystems Sun Solaris
IBM AS/400
IBM MVS
IBM VM
IBM VSE/ESA

Client Platforms

Apple Macintosh
Auspex (via the SunOS client)
Bull DPX/2 300 B.O.S.
Bull DPX/20 (via the AIX client)
Cray UNICOS (via IBM service offering)
Data General DG/UX
Digital OpenVMS (via SSSI's ABC OpenVMS Client for ADSM)
Digital ULTRIX
Digital UNIX
Microsoft DOS
Fujitsu UXP/V and UXP/M (via IBM service offering)
Hewlett-Packard HP-UX

IBM AIX
IBM AS/400
IBM OpenEdition MVS
IBM OS/2 Warp, LanServer, and Warp Server
LINUX - THIS CLIENT IS OFFERED ON AN UNSUPPORTED BASIS ONLY
Microsoft Windows
Microsoft Windows 95 & Windows 98
Microsoft Windows NT DEC Alpha (via the 32-bit DEC Alpha client)
Microsoft Windows NT Intel (via the 32-bit Intel client)
NCR UNIX SVR4 (formerly AT&T GIS SVR4)
NEC EWS-UX/V
Novell NetWare
Pyramid Nile (via the Siemens Nixdorf SINIX)
SCO UNIX 386, Open Desktop, or OpenServer
Sequent PTX
Siemens Nixdorf SINIX 386/486
Siemens Nixdorf SINIX
Siemens Nixdorf SINIX Reliant UNIX
Silicon Graphics IRIX
Sun Microsystems Sun Solaris
Sun Microsystems Sun SunOS
Tandem Guardian (NSK) (via ETI's Backhome for ADSM)

Connect Agents

ADSMConnect Agent for Lotus Notes on AIX and on Microsoft Windows NT Intel
ADSMConnect Agent for Microsoft Exchange Server
ADSMConnect Agent for Microsoft SQL Server
ADSMConnect Agent for Oracle Backup on AIX
ADSMConnect Agent for Oracle Backup on Hewlett-Packard HP-UX, Sun Microsystems Sun Solaris and Microsoft Windows NT Intel

(http://www.storage.ibm.com/software/adsm/adsercli.htm)

Appendix C: Performance Figures for Small and Large Files

NOTE: Performance times depend greatly on network type and traffic.

ADSM for AIX 3.1 on RS/6000 F50 4 way SMP
7133 SSA 9.1 GB Disk Storage Pools
Backup/Restore using Shared Memory Clients

Table 1. Selective Backup, ADSM for AIX Server Version 3 Release 1, RS/6000 F50 4way SMP (PowerPC 604e 166MHz X5DX), 7133 9.1 GB SSA Disk JFS Storage Pools, Shared Memory Local Client (7133 9 GB SSA Disk)

Main Description	Sub Description	Throughput (kb/s, gb/hr)	Thru Ratio	Server CPU Utilization	KB per CPU second
256 mb backup	1 client	5825 kb/s, 20.00 gb/h	100	15%	9708 kb
	2 client	11587 kb/s, 39.78 gb/h	199	31%	9344 kb
	4 client	18396 kb/s, 63.16 gb/h	316	56%	8213 kb
10 mb backup	1 client	5535 kb/s, 19.00 gb/h	100	14%	9884 kb
	2 client	9509 kb/s, 32.65 gb/h	172	27%	8805 kb
	4 client	15170 kb/s, 52.08 gb/h	274	46%	8245 kb
100 kb backup	1 client	2363 kb/s, 8.11 gb/h	100	8%	7384 kb
	2 client	4485 kb/s, 15.40 gb/h	190	17%	6596 kb
	4 client	7271 kb/s, 24.96 gb/h	308	30%	6059 kb
10 kb backup	1 client	1499 kb/s, 5.15 gb/h	100	15%	2498 kb
	2 client	2576 kb/s, 8.84 gb/h	172	29%	2221 kb
	4 client	3021 kb/s, 10.37 gb/h	202	38%	1988 kb
1 kb backup	1 client	220 kb/s, 0.76 gb/h	100	16%	344 kb
	2 client	332 kb/s, 1.14 gb/h	151	29%	286 kb
	4 client	291 kb/s, 1.00 gb/h	132	41%	177 kb

This table is an excerpt from ADSM for AIX Version 3 Release 1 Performance Evaluation Report and shows how times decrease for smaller files.

(http://adsmperf.tucson.ibm.com/adsmax31/adsmax31.html)

Bibliography

Friess, B., Gamet, M., Kadleck, B. & Randall, P. (1998). *Windows NT Backup and Recovery with ADSM*. San Jose: International Business Machines Corporation

Holley, L., Knapp, C. & Paranjape, S. *ADSTAR Distributed Storage Manager (ADSM) for AIX Version 3 Release 1 Performance Evaluation Report*. (1998). San Jose/Tucson: International Business Machines Corporation [Online]. Available: HTTP: http://www.storage.ibm.com/software/adsm/adwhhsm.htm [1998, August 5]

Healy, V. (1997). *IBM User-Centered Design ADSM Competitor Evaluation White Paper*. San Jose: International Business Machines Corporation

Franklin Jr., C.F. (1998, February 15). *The New Storage Explosion*. Solutions Integrator

Script for ADSM Version 3 Enhancements Presentation. (1998). San Jose: International Business Machines Corporation

ADSTAR Distributed Storage Manager (ADSM) - Choosing the Right Storage Management Solution White Paper. (1995). San Jose: International Business Machines Corporation [Online]. Available: HTTP: http://www.storage.ibm.com/software/adsm/adwhchoo.htm

ADSTAR Distributed Storage Manager (ADSM) - Hierarchical Storage Management (HSM) White Paper. (1996). San Jose: International Business Machines Corporation [Online]. Available: HTTP: http://www.storage.ibm.com/software/adsm/adwhhsm.htm

ADSM - Straight Talk with Ron: Column #2 Overview. (1998). San Jose: International Business Machines Corporation [Online]. Available: HTTP: http://www.storage.ibm.com/software/adsm/adtc2p1.htm

ADSM and ADSM Connect Agent Requirements, supported devices, client code and fixes. (1998, September 15). San Jose: International Business Machines Corporation [Online]. Available: HTTP: http://www.storage.ibm.com/software/adsm/adsercli.htm

Story of IBM 1890 - 1938. (1998, September 11). International Business Machines Corporation [Online]. Available: HTTP: http://www.ibm.com/IBM/history/story/

About IBM. (1998, September 11). International Business Machines Corporation [Online]. Available: HTTP: http://www.ibm.com/IBM/

About the IBM Storage Systems Division. (1998, September 11). International Business Machines Corporation [Online]. Available: HTTP: http://www.storage.ibm.com/aboutssd.htm

NOTE: The picture of the Deskstar 16GP hard disk drive on page 2 is from http://www.storage.ibm.com/hardsoft/diskdrdl/prod/deskstar.htm