IBM® TotalStorage® DR550 Performance Measurements

Covering: **Results of benchmark testing on the IBM TotalStorage DR550**

Updated as of February 14, 2005

IBM Storage Products
http://www.storage.ibm.com

IBM TotalStorage DR550
http://www.storage.ibm.com/disk/dr

Table of Contents

Disclaimer

February 14, 2005

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of the intellectual property rights of IBM may be used instead of the IBM product, program, or service.

While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The performance information contained in this document was derived under specific operating and environmental conditions. The results obtained by any party implementing the products and/or services described in this document will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected in any implementation of such products and/or services. Accordingly, IBM does not provide any representations, assurances, guarantees, or warranties regarding performance.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, INTEROPERABILITY, OR NON-INFRINGEMENT.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time. The information contained in this document is current as of the initial date of publication only, and IBM shall have no responsibility to update such information.

Trademarks

IBM, the IBM logo, AIX, POWER5, RS6000, Tivoli, and TotalStorage are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Other company product and service names may be trademarks or service marks of their respective owners.

Notices

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, New York 10594, U.S.A.

Introduction

The IBM TotalStorage DR550 offering designed to help companies manage (archive and retrieve) data assets in a non erasable and non rewritable format. While not all data has a requirement to be stored in non erasable and non rewritable formats, new industry and government regulations are increasing the amount and type of data to be stored in this type of format. Data Retention offerings such as the DR550 are designed to archive and manage data that needs to be retained for a specific period of time, and while retained, can not be modified or deleted.

The DR550 includes IBM hardware and software, specifically IBM's POWER5 servers with AIX, Tivoli Storage Manager for Data Retention, and TotalStorage DS4100 Disk System. The offering is an integrated offering where the hardware has been installed, cabled and tested in a rack and the software has been pre-configured. System administrators are required to define the TSM management classes associated with the retention periods for their business needs. When combined with a content management application, data can be managed throughout its lifecycle, including automatic deletion at the right time. The combination of hardware and software in the DR550 creates an environment where modification or erasure of data is prevented until the retention period is reached.

Objective of this paper

This paper has been written to provide performance and capacity planning information for the DR550.

Configuration

A 7 terabyte DR550 (detailed configuration is included in the appendix) was used for the measurements as set forth in this paper. It was directly attached through a single gigabit Ethernet connection to an IBM RS6000 Enterprise Server Model M80 (running AIX 5.2.0).

Environments

An application program, using the standard Tivoli Storage Manager (TSM) Application Program Interface (API), was used to create a variety of TSM workloads. This program was designed to archive and retrieve sufficient objects to characterize the operational performance characteristics of the DR550. Object sizes used were 8K, 16K, 32K, 64K, 128K, 256K, 512K, and 1024K. The application ran on the M80.

The TSM API allows for the archiving and retrieving of multiple objects in a single session. Thus the session setup time required to establish network communication, check parameter lists, etc. need only be done once, rather than per object. This enables the TSM client to archive all of the modified files for a particular client in a single session. The performance results were characterized at 1, 2, 10, 25, 100, 250, and 500 objects per session for multiple object sizes using a single client.

Also, the 250 objects per session environment was exercised for 2 and 4 clients making requests in parallel.

Measurements

Figure 1 below shows the object archive rate for various size objects while varying the number of objects processed per session. A session is the work associated with one interaction between the TSM client and the TSM server.

Archive Performance

Figure 1 shows that the DR550 objects archived per second rate can vary from approximately 5 to over 800 based on the object size and the number of objects per session. Figure 9 on page 13 includes a table which shows the actual values.

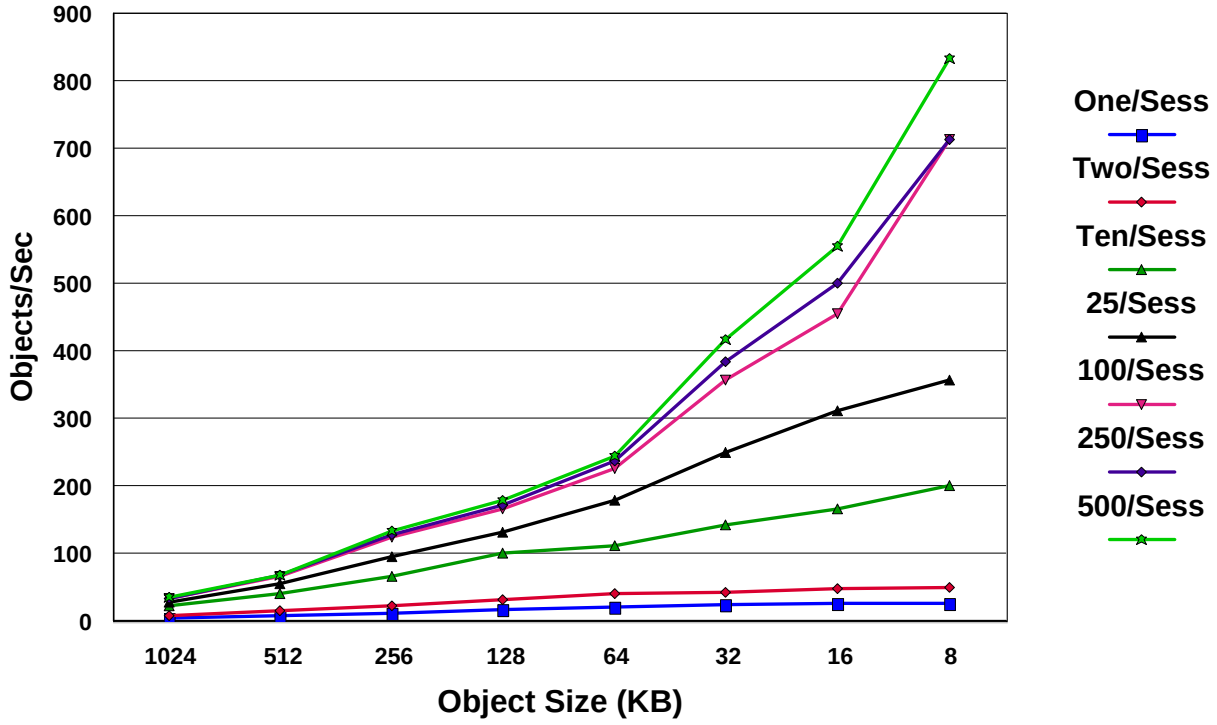

Figure 1 – Archive objects/second

The following observations have been made:

- There is a large difference in the number of objects per second which can be transferred depending on the size of the objects. That is, knowing the average size of the objects you process is tremendously important to your capacity planning effort.
- Knowing the average size of the objects being processed is critical to capacity planning since the objects per second transferred depends primarily on the size of the objects. Thus, it is critical to understand if the content management application being used will transfer multiple objects within a single session.

The following is a capacity planning example for an e-mail application.

- 30,000 users each retaining 30 messages a day on average. (Users also delete another 50 messages which are the electronic equivalent of “junk mail”).
- The average messages size is about 20K.
- We know that the e-mail application will process multiple objects in a single session.

With archiving done by user (each user would be one session), the closest data point (to the 30 messages per day) is 25 messages per session. A bit conservative, but we're not all that sure about the 30 messages per employee, so conservative is good. Following the 25 objects per session line on the graph in figure 1 on page 5, we estimate the DR550 can do 300 objects per second if the average object size is a bit over 16K. Since we are going to archive 30 messages for each of the 30,000 people using the system, we estimate the DR550 will archive 900,000 objects in 3,000 seconds, or 50 minutes.

Retrieve performance

At times, there will be a request to retrieve an old object. We are interested in how long it will take to retrieve this e-mail. This is a single request for a single e-mail. Refer to Figure 2 below using the One/Sess line. Figure 10 on page 14 has more granular data. We find that 13.7 objects a second were transferred when the object size was 16 KB and 14.1 objects per second when the objects were 32 KB. Thus, the object size was not a significant performance parameter. If you retrieve 14 objects a second without overlap, the average time per object is 71.4 milliseconds (MS).

Data rate

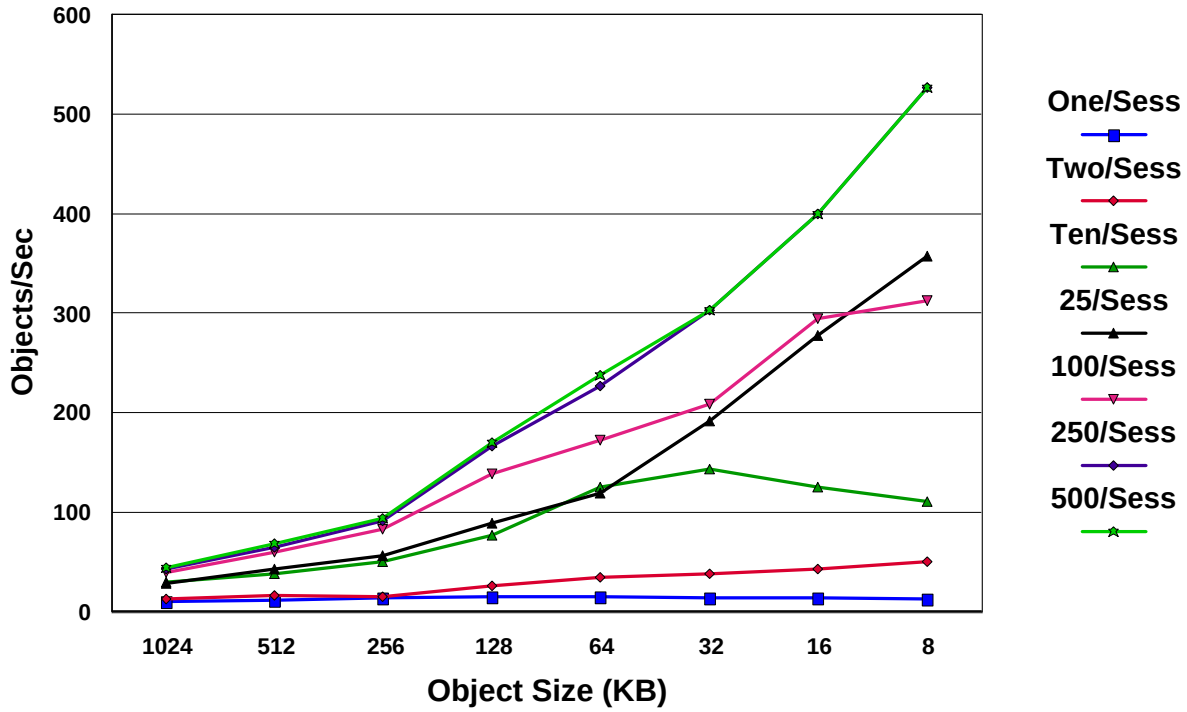

Figure 2 – Retrieve objects/second

If the application uses large objects, we are much more likely to prefer looking at the data transfer rate sustained.

Figure 3 shows there is very little change in transfer rate when we ran between 100 and 500 objects per session. While the transaction rates were lower for large block sizes, the data transfer rates were much higher.

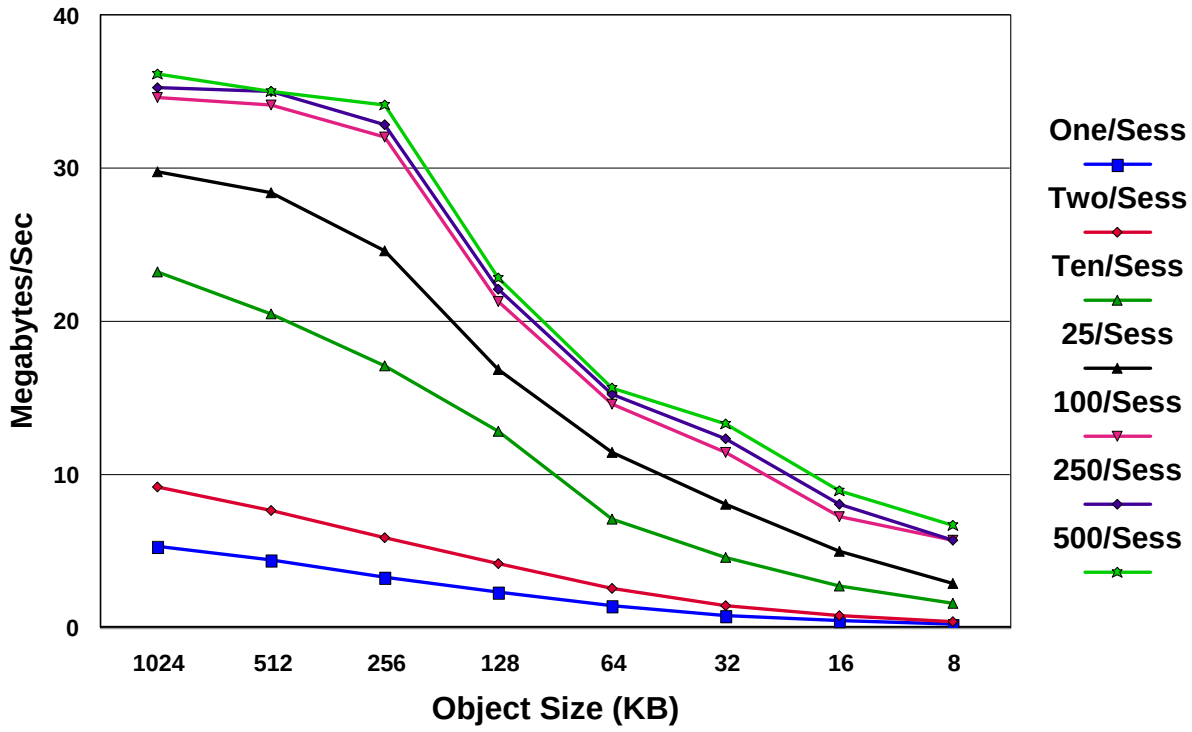

Figure 3 – Archive Megabytes/second

If you are archiving data from digital cameras and each camera generates 1 GB of data per hour (.28 MB/sec), you could archive data from 100 cameras and be well within the guidelines (35 MB/sec with 100 objects/session).

The data rate chart in Figure 3 is a more useful capacity planning tool for applications with very large object sizes.

Retrieve transfer rates are shown below in figure 4. These might be useful for large object applications.

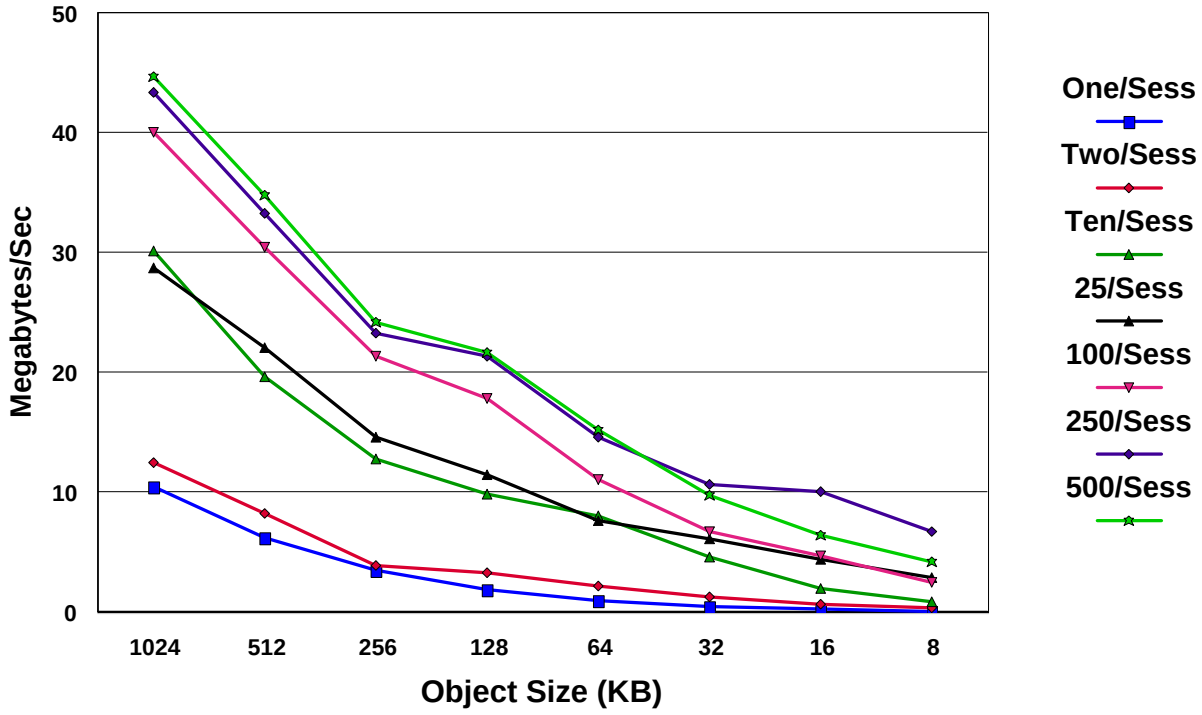

Figure 4 – Retrieve Megabytes/second

For example, if a movie took a gigabyte of storage, it would take 100 seconds to retrieve at 10 megabytes a second. Naturally, the beginning of the movie could start playing a lot faster than that, but, handling that size without any parallel operations took our DR550 about 100 seconds in total.

If the clients supporting 2 different video editors requested a movie at the same time, it would take about 200 seconds to completely retrieve 2 movies. There would be instances where one client was ignoring a resource (TSM server, Ethernet connection, DS4100 control unit, etc.) and the other client could use that resource. In the lab, we ran a single client in the interest of understanding the resources required. In most real uses we would expect multiple clients to be using the TSM server.

Parallel requests

Below are the results when archiving from one, two, and four clients. They all used the same TSM client in the RS6000 model M80 and the single Gigabit Ethernet connection.

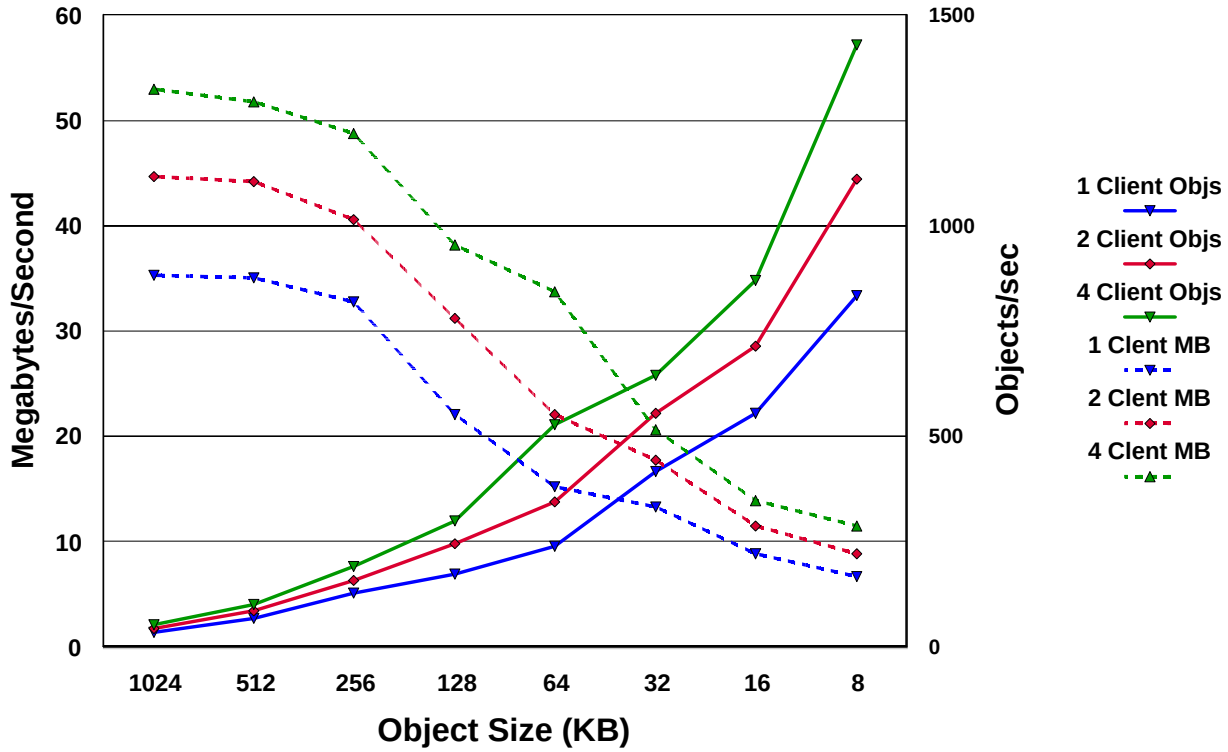

Figure 5 – Concurrent processing

For the megabyte object size, there was a 26% improvement going from 1 to 2 clients, 50% improvement going from one to four clients. The eight kilobyte objects improved (objects per second) by 33% moving from 1 to 2 and 21% moving from 1 to 3.

Resource usage

Why didn't we get twice as much work done when there were twice as many clients? Because some resource is running near saturation and cannot run twice as much work no matter how many clients make requests.

DR550 Server utilization

Figure 6 shows the processor utilization on the DR550 server at different times during the run. If we had an objective of archiving more objects we would be very interested in making available more processors or faster processors.

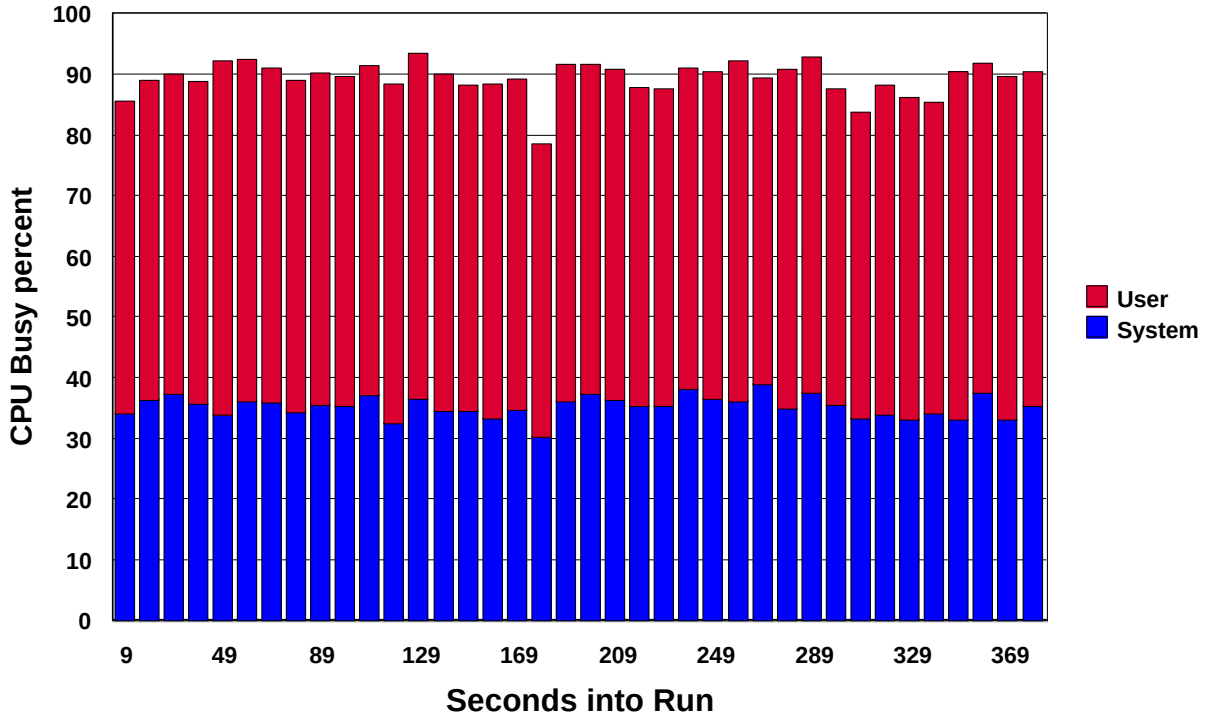

Figure 6 – Server Utilization

Quite possibly, some other resource is also near saturation. It could be the Ethernet connection, we saw in figure 5 on page 9 that the maximum we could drive it using 4 concurrent clients was 52.92 MB/second. The connection is rated at 100 MB/second.

Well, maybe we are having difficulties with the TSM data base. No, it doesn't seem likely as the percent busy for the hdisks (as shown in figure 7) on which the TSM Database and logs were located were occasionally a very low number, rather than zero. The peak value was 2.2% busy across all 3 hdisks.

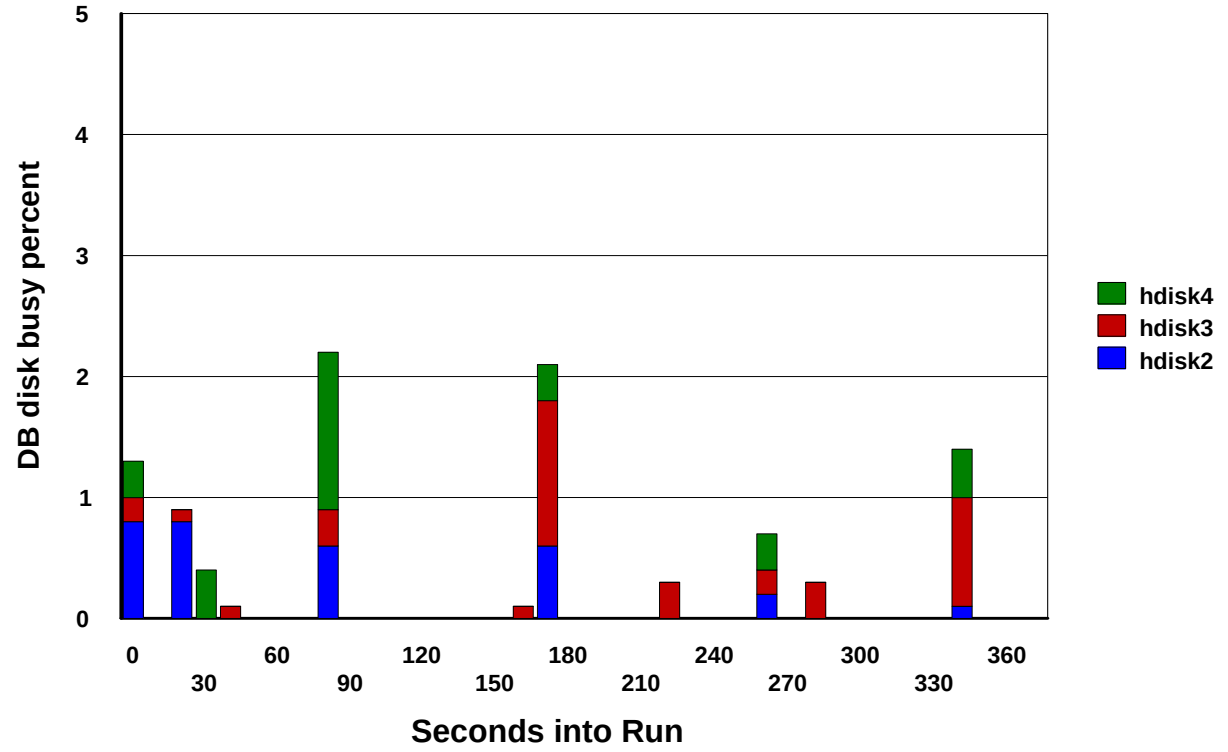

Figure 7 – DB disk busy percent

The hdisks used for object storage were much more busy as shown in figure 8.

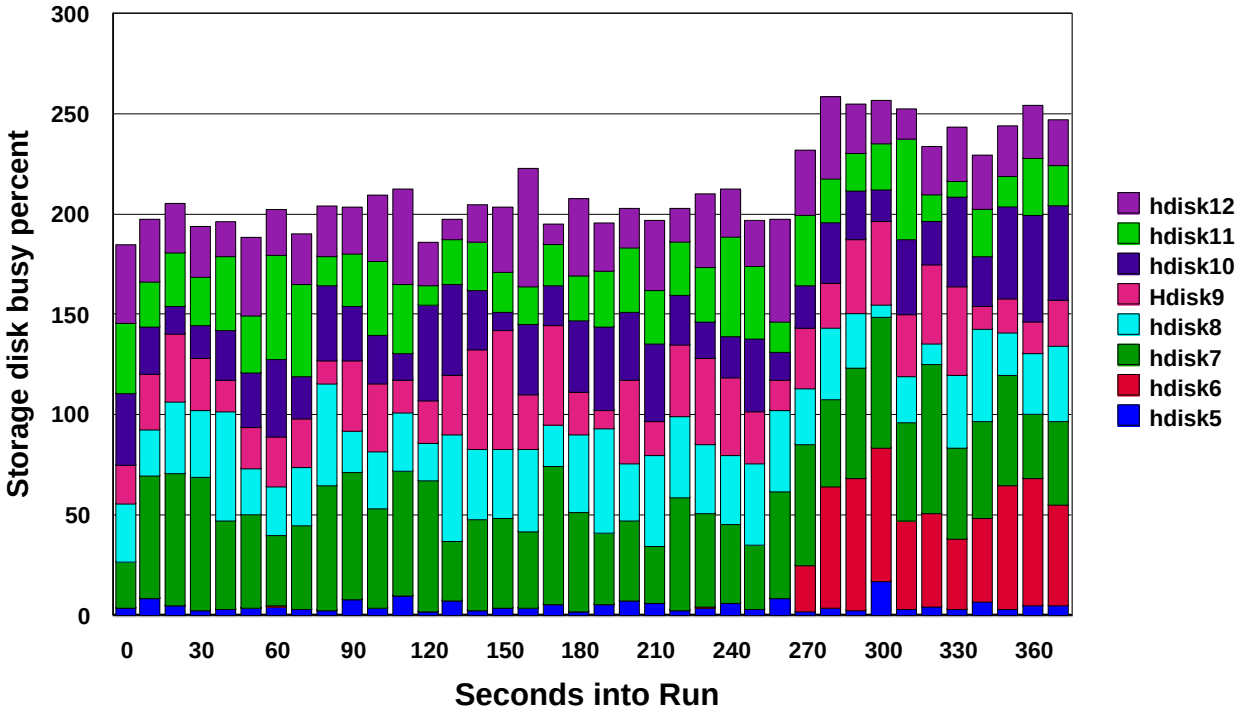

Figure 8 – HDISK utilization

While the utilization was higher than that shown for the database, it was not really a great cause of concern. Total for all activity peaked at about 250%. That is less than 32% for each of the 8 hdisks, a little over 10% for each of the SATA drives actively used. It would seem we could double the processor cycles available and go quite some ways before Ethernet connections or DS4000 equipment becomes a limiting factor.

Conclusion

DR550 can offer excellent performance, especially when processing sessions with more than a single object. Planning for the appropriate DR550 configuration should be done as part of the overall project. The DR550 has been tuned to provide balanced performance for both small and larger disk capacities. The number of DR550s required should be based on the number of objects to be archived and retrieved and the ability of the content management application to support multiple objects per session.

Appendix

Tabular data

These tables are provided for increased accuracy over the corresponding graphs.

Objects per second

Archive

Object Size	One/ Sess	Two/ Sess	Ten/ Sess	25/ Sess	100/ Sess	250/ Sess	500/ Sess
1024	Sess	Two/	22.73	29.07	33.78	34.48	35.34
512	Sess	Ten/	40.00	55.56	66.67	68.49	68.49
256	Sess	25/	66.67	96.15	125.00	128.21	133.33
128	Sess	100/	100.00	131.58	166.67	172.41	178.57
64	Sess	250/	111.11	178.57	227.27	238.10	243.90
32	Sess	500/	142.86	250.00	357.14	384.62	416.67
16	Sess	47.62	166.67	312.50	454.55	500.00	555.56
8	27.03	50.00	200.00	357.14	714.29	714.29	833.33

Retrieve

Object Size	One/ Sess	Two/ Sess	Ten/ Sess	25/ Sess	100/ Sess	250/ Sess	500/ Sess
1024	10.2	12.20	29.41	28.09	39.06	42.37	43.67
512	12.05	16.13	38.46	43.10	59.52	64.94	68.03
256	13.51	15.15	50.00	56.82	83.33	90.91	94.34
128	14.71	25.64	76.92	89.29	138.89	166.67	169.49
64	15.15	34.48	125.00	119.05	172.41	227.27	238.10
32	14.08	38.46	142.86	192.31	208.33	303.03	303.03
16	13.70	43.48	125.00	277.78	294.12	400.00	400.00
8	12.99	50.00	111.11	357.14	312.50	526.32	526.32

Figure 9 – Archive/Retrieve objects/second

Throughput (MB per second)

Archive

Object Size	One/ Sess	Two/ Sess	Ten/ Sess	25/ Sess	100/ Sess	250/ Sess	500/ Sess
1024	5.31	9.14	23.27	29.77	34.59	35.31	36.18
512	4.38	7.64	20.48	28.44	34.13	35.07	35.07
256	3.24	5.82	17.07	24.62	32.00	32.82	34.13
128	2.29	4.13	12.80	16.84	21.33	22.07	22.86
64	1.39	2.56	7.11	11.43	14.55	15.24	15.61
32	0.78	1.39	4.57	8.00	11.43	12.31	13.33
16	0.41	0.76	2.67	5.00	7.27	8.00	8.89
8	0.22	0.40	1.60	2.86	5.71	5.71	6.67

Retrieve

Object Size	One/ Sess	Two/ Sess	Ten/ Sess	25/ Sess	100/ Sess	250/ Sess	500/ Sess
1024	10.45	12.49	30.12	28.76	40.00	43.39	44.72
512	6.17	8.26	19.69	22.07	30.48	33.25	34.83
256	3.46	3.88	12.8	14.55	21.33	23.27	24.15
128	1.88	3.28	9.85	11.43	17.78	21.33	21.69
64	0.97	2.21	8.00	7.62	11.03	14.55	15.24
32	0.45	1.23	4.57	6.15	6.67	10.67	9.70
16	0.22	0.70	2.00	4.44	4.71	10.00	6.40
8	0.10	0.40	0.89	2.86	2.50	6.67	4.21

Figure 10 – Archive/Retrieve data rates

DR550 Configuration

The following products were included in the 7 Terabyte DR550 that was used in these measurements.

Hardware

POWER5 520 (9111-520)		1
0265	AIX Partition Specify	1
0706	Integrate with IBM TotalStorage DR550 (Single node)	1
2849	POWER GXT135P Graphics Accelerator With Digital Support	1
3273	36.4 GB 10,000 RPM Ultra320 SCSI Disk Drive Assembly	2
4444	1024 MB (4x256 MB) DIMMs, 208-pin, 266 MHz DDR SDRAM	1
5005	Software Preinstall	1
5158	AC Power Supply, 850 W	1
5231	1-way 1.5 GHz POWER5 Processor Card, No L3 Cache	1
5706	IBM 2-Port 10/100/1000 Base-TX Ethernet PCI-X Adapter	1
5716	2 Gigabit Fibre Channel PCI-X Adapter	3
5751	IBM 4.7 GB IDE Slimline DVD-RAM Drive	1
6458	Power Cable -- Drawer to IBM PDU, 14-foot, 250V. 10A	1
6574	Ultra320 SCSI 4-Pack	1
7160	IBM Rack-mount Drawer Rail Kit	1
7877	Media Back plane Card	1
7918	IBM Rack-mount Drawer Bezel and Hardware	1
8462	Zero-priced Value Pak Processor Entitlement for # 5231	1
9300	Language Group: US English	1
DS4100 Midrange Disk System		1
0707	Integrate with IBM TotalStorage DR550	1
2210	(19K1271) Short Wave SFP GBIC (LC)	4
4602	(90P1350) SATA 250 GB/7200 Disk Drive Module	14
5605	(19K1248) LC-LC 5M Fibre Optic Cable	2
7711	(25R0063) DS4100 AIX Host Kit	1
DS4000 EXP100 Storage Expansion Unit (1722-10U)		1
0707	Integrate with IBM TotalStorage DR550	1
2210	(19K1271) 2Gb Fibre Channel Short Wave GBIC	2
4602	(90P1350) SATA 250 GB/7200 Drive Module	14
5601	(19K1247) 1m 50u Fiber Optic Cable (LC-LC)	2
9001	Attach to DS4100 (1724-100)	1
IBM TotalStorage SAN Switch - 8 Port (2005-H08)		1
0706	Data Retention DR550 single node-Integration	1
2214	2 Gbps SW SFP Transceiver - 4 Pack	1
5605	Fibre Cable LC/LC 5m multimode	3
7310-CR2 Rack-mounted Hardware Management Console (7310-CR2)		1
0706	Integrate with IBM TotalStorage DR550 (Single node)	1
0961	Hardware Management Console for POWER5 Lic Machine Code	1
4651	Rack Indicator, Rack #1	1
6458	Power Cord (14 ft) 250V/14A, IEC320/C13, IEC320/C14	1
7801	Ethernet Cable, 6M, Hardware Management Console System Unit	1

9300	Language Group: US English	1
IBM Rack-Mounted Flat Panel Console Kit (7316-TF3)		1
0706	Integrate with IBM TotalStorage DR550 (Single node)	1
4202	Keyboard/Video/Mouse (LCM) Switch	1
4242	6 Foot Extender Cable for Displays	1
4269	USB Conversion Option	1
6350	Travel Keyboard, US English	1
9300	Language Group Specify - US English	1
9911	Power Cord (4M) All (Standard Cord)	1
IBM RS/6000 Rack Model T00 (7014-T00)		1
0229	Rack Content Specify: 9111/520 - 4 EIA	1
0233	Rack Content Specify: 7316-TF3 - 1 EIA	1
0234	Rack Content Specify: 1710/10U - 3 EIA	1
0247	Rack Content Specify: 7310/CR2 - 1 EIA	1
0248	Rack Content Specify: 1 EIA	1
0254	Rack Content Specify: 1724/100 - 3 EIA	1
0706	DR550 Integration (single node)	1
6068	Front Door For 1.8 Meter Rack (Flat Black)	1
6098	Side Panel For 1.8 or 2.0 Meter Rack (Black)	2
6580	Optional Rack Security Kit	1
7176	Single Phase PDU 30A L6-30 cable	1
9176	Single Phase Power 30A Specify	1
9300	US English Nomenclature	1
9800	United States/Canada Power Cord	1

Software

Tivoli

	IBM Tivoli Storage Manager for Data Retention (5608-DR3)	1
	IBM Tivoli Storage Manager Extended Edition (5608-ISX)	1
AIX System Software (5692-A5L)		1
0947	MEDIA 5765-E62 AIX V5.2	1
0949	Expansion pack	1
0965	Aix 5.2 Update CD	1
0975	Microcode Upd Files and Disc Tool v1.1 CD	1
1004	CD-ROM Process Charge	1
1432	Preinstall Tivoli Security Agent	1
2924	English Language	1
3410	CD-ROM	1
5005	Preinstall	1
5924	English U/L SBCS Secondary Language	1
AIX V5 (5765-E62)		1
0009	Value Pak Per Processor E5 AIX V5.2	1

