IBM® System Storage™ DS5300
Performance Results in IBM i™ Power Systems Environment

Performance White Paper

December 2008

Wesley Varela
IBM i™ Performance

Systems and Technology Group
International Business Machines Corporation

IBM System Storage DS5300 with seven EXP5000 expansion enclosures in a rack

Table of Contents

Introduction

With more complex customer storage requirements, Storage Area Network (SAN) solutions are becoming the de facto standard. IBM i™ Performance would like to present the DS5300™ Disk System performance measurements. The DS5300™ is an excellent Mid to High range storage system that increases performance and capacity relative to the high end DS4000 series by leveraging the latest software and hardware technology. The DS5300™ is perfectly suited to leverage the benefits inherent in the 4 Gbps Fiber Channel technology with sixteen 4 Gbps host channels to quickly transfer data, as well as, the 16 GB of physical cache memory (8 GB per controller). The DS5300™ can access a total of sixteen EXP5000 or EXP810 expansion drawers for a total of 256 attached disk drives. Typical high throughput and data-intensive applications that would benefit from the DS5300™ include (but not limited to) V.O.D (Video On Demand), Medical Imaging, Data mining/warehousing, advanced replication services and, of course, Online Transaction Processing (OLTP).

In an Effort to show how the DS5300™ compares to an IBM i internal storage solution we have completed several tests using CPW (Commercial Processing Workload) and various Save/Restore workloads using RAID5, RAID6 and RAID10 to demonstrate the particular characteristics of the DS5300™ as compared to a completely internal solution. While there are a great many reasons to consider an external storage solution, we highly recommend reviewing several of the documents available online and contacting your IBM Representative to help you find the best external storage solution for your business requirements.

IBM System Storage Product Guide
http://www-03.ibm.com/systems/resources/systems_storage_resource_pguide_prodguidedisk.pdf

IBM System Storage DS5000 series
http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=PM&subtype=SP&appname=STG_TS_USEN&htmlfid=TSD03061USEN&attachment=TSD03061USEN.PDF

Test Environment and System Configurations

All measurements where taken using CPW* and Save/Restore. CPW is an internal tool used for Online Transaction Processing (OLTP) measurements.

- *Commercial Processing Workload (CPW).*
  This workload is characterized by many jobs running brief database transactions in an environment that is dominated by IBM system code performing these database operations. CPW is used by IBM for product evaluation of different IBM i models where the primary application is oriented to traditional commercial business uses (order entry, payroll, billing, etc.) using commitment control. Although the CPW workload is typically used to represent the relative performance of a complex of processors, it does a significant number of read and write DASD accesses capable of saturating the I/O subsystem making CPW particularly useful in disk characterization.
- *Save and Restore Workloads.* These workloads include:
  - *Large Database File.* This is a single database file with one member. It is 64 gigabyte(GB) in size.
  - *User Mix.* This is a mix of objects including source files, database files, programs, command objects, data areas, menus, query definitions, etc. It is 12GB in size with approximately 49,500 objects.

IBM i™ Power System configuration used in these tests**:

- Seven processors of an IBM i 8 way 9117-MMA P6 570.
- Both Internal and External configurations used the same 8 disk system ASP RAID5 which was hosted in the DASD.
- 128GB main store where used on both internal and external partitions
  - Internal Configuration:
    - 96 Disks
      - 72 - Database ASP Mirrored, RAID5 and RAID6
      - 24 - Journal ASP unprotected
    - Six 571B storage adapter cards (90 MB cache per card)
      - 4 for Database ASP, 2 for Journal
  - DS5300™ 1818-53A Configuration***:
    - 96 DDMs
      - 12 LUNS (72 DDM) – Database ASP Mirrored, RAID5 and RAID6
      - 4 LUNS (24 DDM) – Journal ASP unprotected
    - Two 576B 4Gb FC adapter
      - Each port on both cards where connected to different controllers
      - Each line was connected to different cards in the controller.
      - (2) DS5300™ Host Groups: (6) Database LUNS and (2) Journal LUNS comprised each Host group
    - VIOS was running on 1 CPU with 2 GB main store.

- The DS5300™ 1818-53A has 16 GB cache (8 GB per controller)
- 6 EXP810 Enclosures where used.
- Each LUN was constructed using one disk from each expansion as illustrated in figure 1.

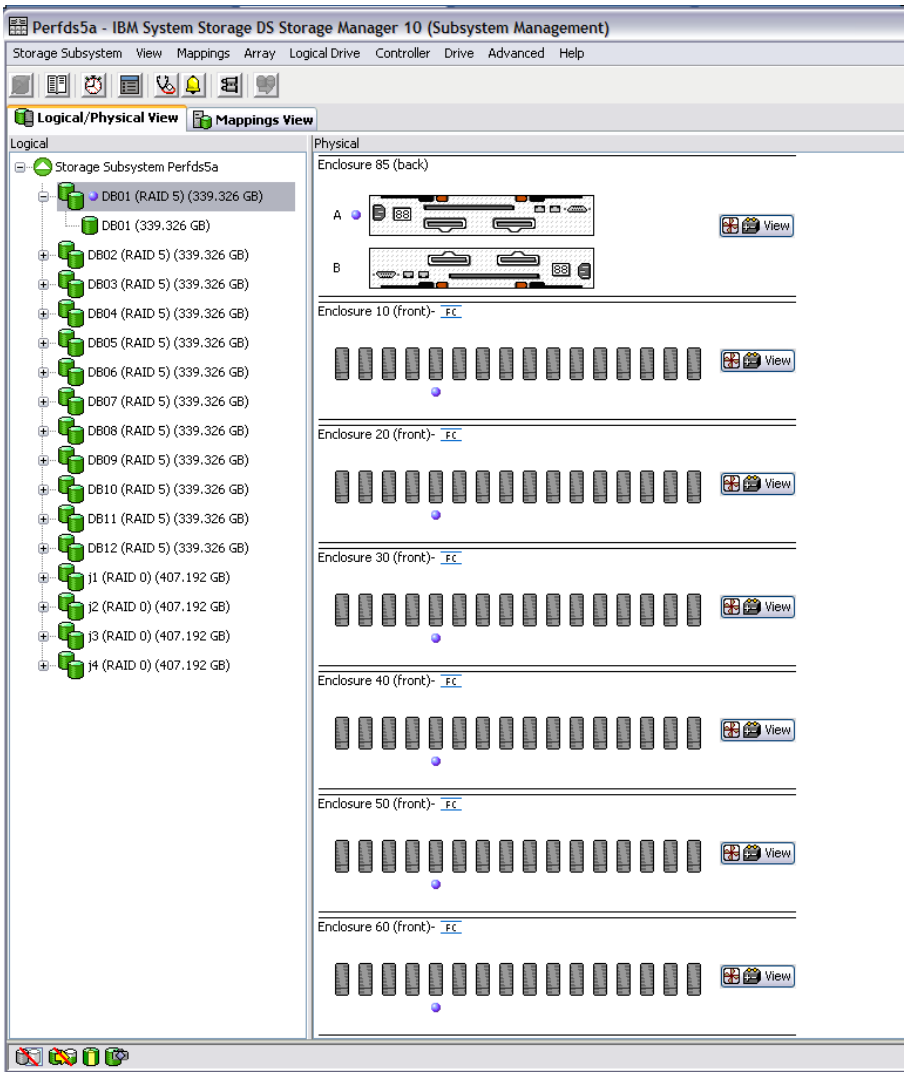

Figure 1

- IBM i 6.1
- Type 4328 (146GB, 15K RPM) internal disks attached to the storage adapters

* During both internal and external evaluations, no other system activity was present.
** Tests were designed to compare equal number of physical drives not equal capacity.
*** External configuration hosted the Journal ASP and the Database ASP. System ASP was on the IBM i Power System. The DS5300™ was configured using Storage Manager V10.30.

Performance Results

CPW - User Series

Using CPW we are able to test the system configuration using a wide array of methodologies. One such method is a "*User Series*" test, in this test a specific number of "*users*" are released at discrete intervals. During each interval we can evaluate the I/O subsystem to monitor disk characteristics as a whole, per ASP, per controller even per LUN, as well as over all application response time. For our internal and external configuration tests we started at 16K users and released an additional 16K users at each time interval to a maximum of 176K users with each user attempting to perform 1.22 transactions per minute in each time interval.

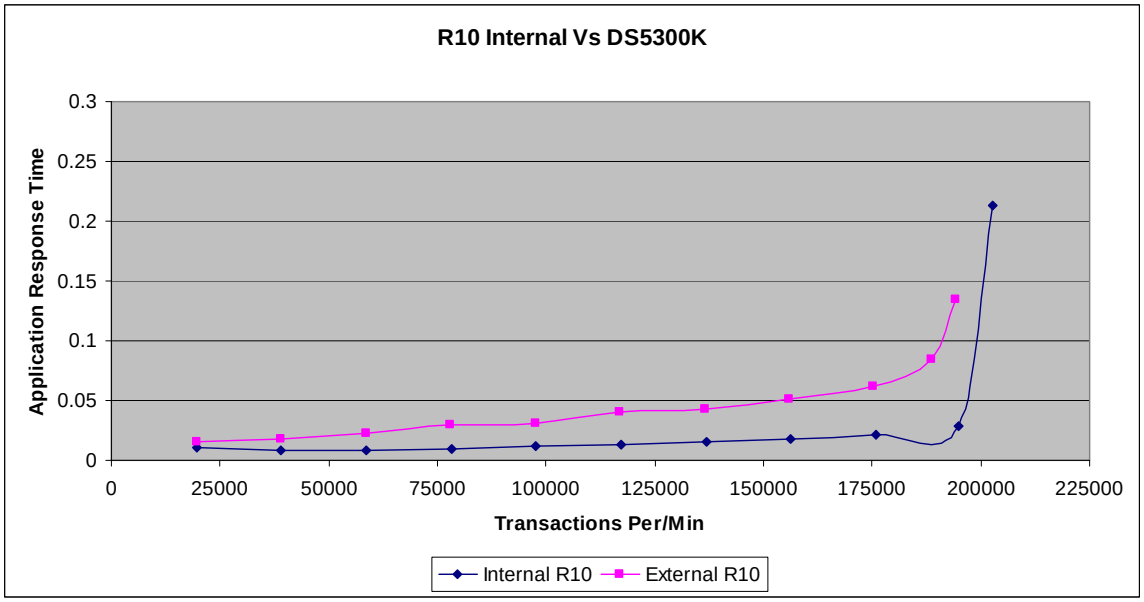

Figure2 Application Response Time

As our tests were performed with the CPW *"user series"* method, each point on the charted line in Figure 2 represents an additional 16K users being added to the workload. Figure 2 shows the disk response time for internal and external DDM as the number of operations increase to the database ASP. If we examine the External RAID10 (mirrored) line and the Internal RAID10 (mirrored) line we will see a slight increase in the (average) disk response time as the number of operations per second increase, which is not unexpected; management overhead on any external storage solution will be present. However it is important to note that both mirrored runs had an impressive Application Response Time of 1 millisecond at the high end of the workload. This effectively mitigates the trade off between the tremendous flexibility of an external storage solution (FlashCopy, VolumeCopy and Enhanced Remote Mirroring) and the increased latency over internal disks. Figures 3 examine the database ASP and Journal ASP separately.

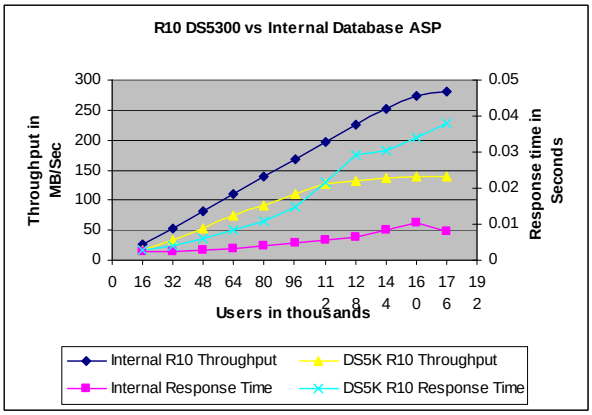

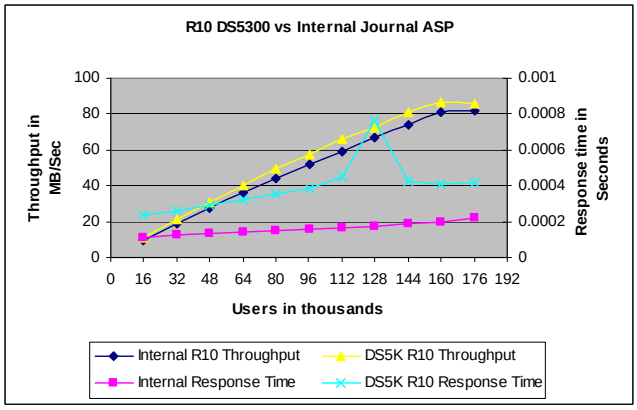

Figure3 Database and Journal ASP Response Times

It is important to point out that Journal LUNS shared a host group connection with Database LUNS, with an increase in response time in a single Host Group can affect but the Journal and the Database ASP response times.

Figure 5 examines the RAID5 and RAID6 solutions on the DS5300, while figures 6 and 7 examine the individual ASP's. As you can see, the DS5300™ has advanced caching algorithms and dedicated CPU's that allow it to excel. The DS5300™'s internal cache handily overpowers the internal solutions adaptor cards for the RAID5 and RAID6 workloads.

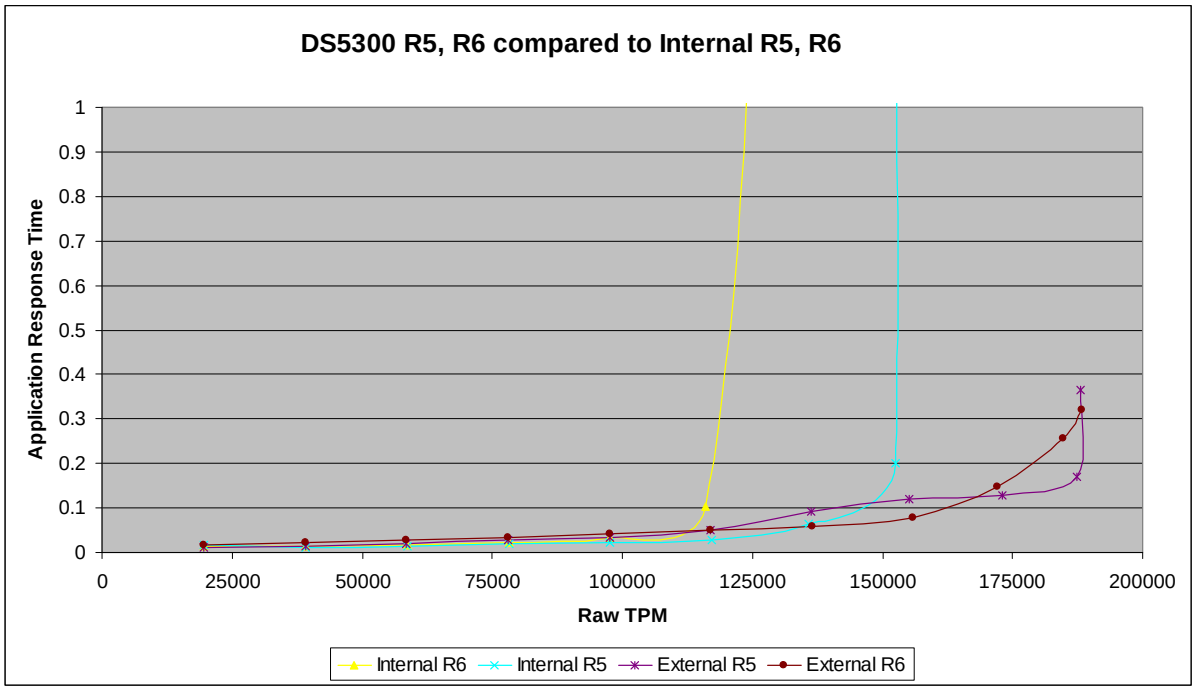

Figure 5 R5 and R6 comparisons

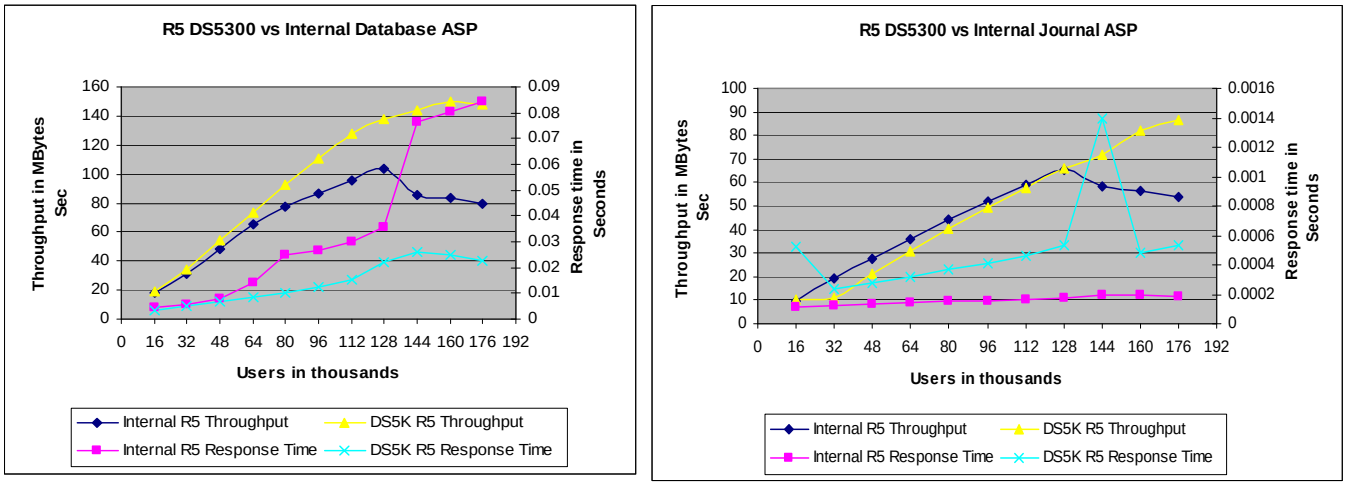

Figure 6 R5 Database and Journal ASP Views

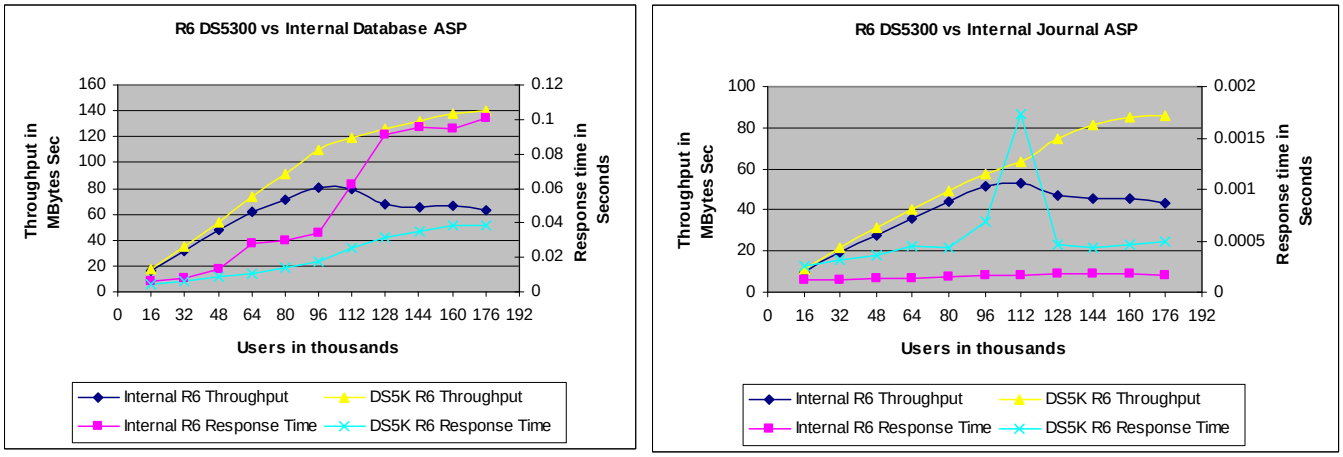

Figure 7 R6 Database and Journal ASP Views

Recall that CPW is an OLTP (Online Transaction Processing) workload. While latency on individual disk response time increases slightly in a DS5300™ the application response time and transactions per/min that are the hallmarks in this category of testing, were very impressive.

Application response time is a weighted average of all transaction response times in a given time interval. Transactions per/min indicate the raw number of accumulated transactions in that interval. Figure 2 through 3 clearly indicate that the DS5300™ performs impressively as compared to internal storage using the CPW workload.

Save and Restore

Save and Restore Workloads

It is important to recall that Save and Restore performance is both data and system specific. Results may vary based on the data, platform and configuration used in the testing. When comparing individual results to the results achieved in this document, it is important to keep these consideration in mind.

Note that these save and restore measurements were all done to save files. To determine save and restore rates to other various devices see the *IBM Power Systems Performance Capabilities Reference IBM i operating system Version 6.1* in the references section of this document. Chapter 15 on Save Restore Performance focuses on save and restore rates and factors that affect them. Chapter 15.12 specifically addresses virtual tape, which was the implementation in this document.

Configuration

Measurements were conducted to test the performance of DS5300 external storage RAID 5, RAID 6 and RAID 10 versus internal storage RAID 5, RAID 6 and RAID 10 (mirroring) while running two specific performance save/restore workloads, large database file *(single database where all members are at least 4 GB in size)* and user mix *(single library made up of a combination of source files, database files, programs, command objects, etc)*. The external or internal disks under test were placed in ASP 2 which contained the save file. ASP 9 contained the workloads. Data was transferred from ASP 2 to ASP 9 for the save and from ASP 9 to ASP2 for the restore.

ASP 9 is a virtual tape consisting of 72 type 4328 (146GB, 15K RPM) internal disks attached to (2) 571F I/O Adapters were used to insure any limiting factors would not be related to the configuration of the virtual tape and adequately test the internal and external storage solutions.

Here are the results of the save and restore runs:

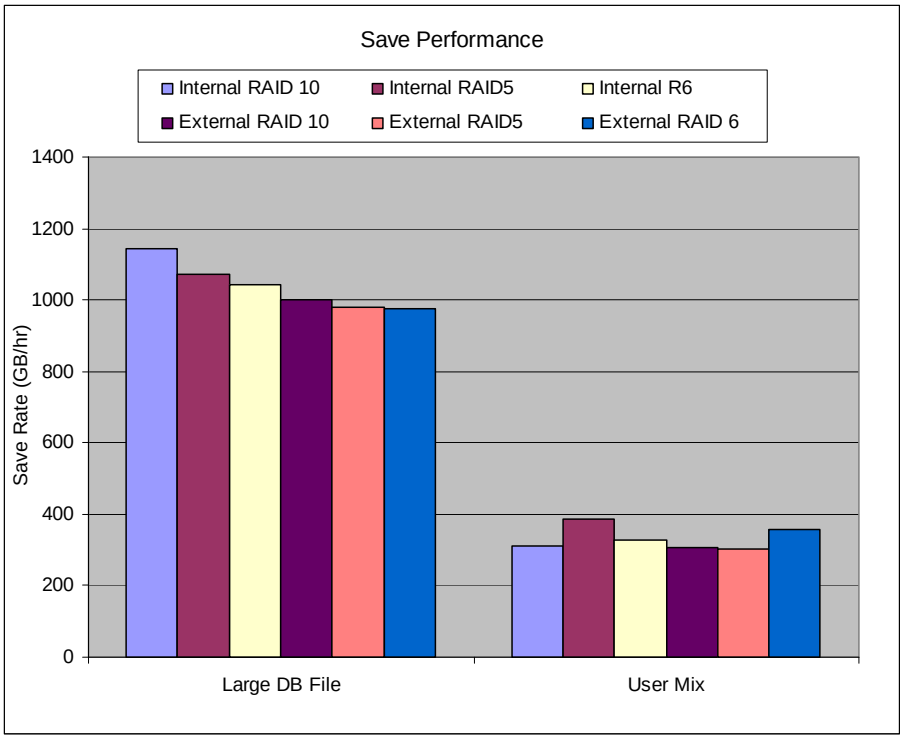

Figure 4 Save performance measurements

The Save portion (figure 4) of the workload shows that the performance is a function of disk response time in ASP 9 as well as the utilization of the 4 GB Fiber Chanel connection between the DS5300™ and our IBM i platform. The external storage solution performed exceptionally well, approaching 1000 GB/hr save rate on all large database file workloads, clearly placing it in the same equivalence class as the internal storage solution. Note that the tests were designed to compare **equal physical hardware** and not equal capacity between RAID and mirroring, which for mirroring, effectively halves the number of arms, creating the drop in the save rate for the internal mirroring configuration compared to the internal RAID 5 configuration.

The Restore portion of the workload effectively leverages the 16 GB of cache to significantly outperform the internal storage solution (figure 5). In the workload that tests the restore of large database files, the DS5300™ almost doubles the restore rate of the internal storage solution and performs equivalently on the user mix files.

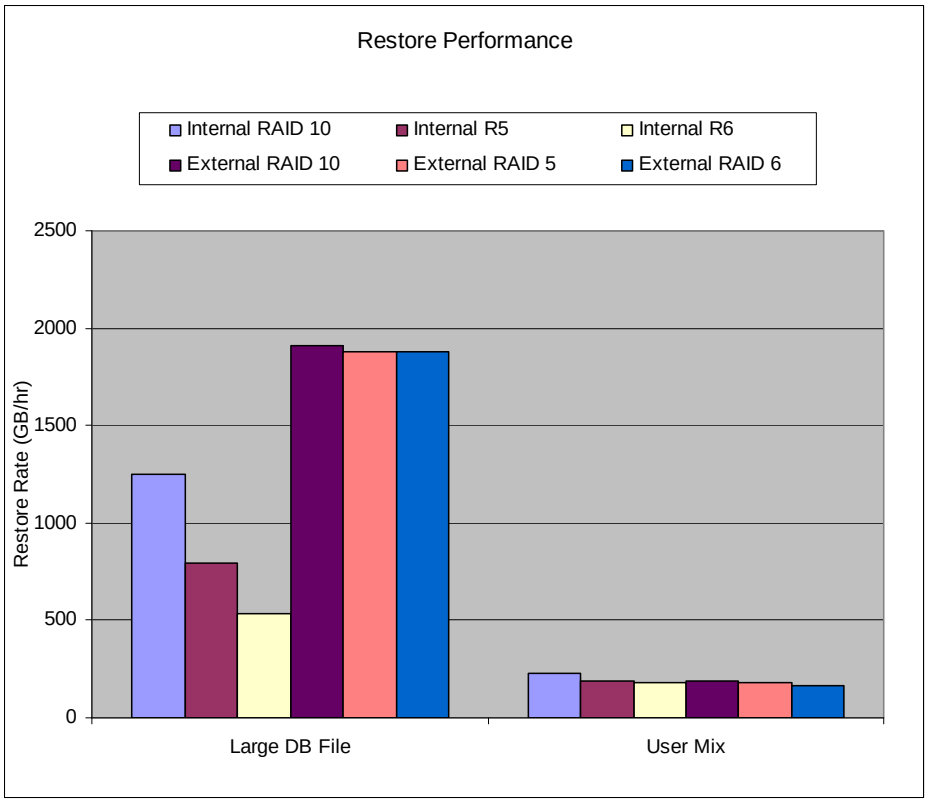

Figure 5 Restore performance measurements

Conclusions

Incorporating a DS5300™ into any environment has its advantages and disadvantages. The DS5300™ up to offers sixteen 4 Gbps Fiber Channel interfaces for high-performance host connectivity that will deliver high throughput and sustained bandwidth for high performance applications. Additionally, connecting to a Storage Area Network (SAN) over a high speed network can offer tremendous flexibility and reliability.

While additional latency exists in the storage management of any SAN product the DS5300™ offers impressive performance results compared to a completely internal solution. Proper planning and sizing considerations are essential in effectively deploying any SAN solution. Please consult Chapter 5 of "Sizing external storage for i5/OS" in the *iSeries and IBM TotalStorage: A Guide to Implementing External Disk on IBM eServer i5* Redbook(s) for more information. Both Redbooks are linked in the reference section of this paper.

Disclaimer – System Performance

Performance results set forth in this document are based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience may vary depending upon considerations, including, but not limited to, the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance characteristics or improvements equivalent to the ratios stated here.

Notices and Trademarks

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

References

IBM Power Systems Performance Capabilities Reference IBM i operating system Version 6.1 Performance Management, published October 2008,
http://www-03.ibm.com/systems/resources/pcrm_oct2008.pdf

DS5000 Disk Storage Subsystem Architecture, Implementation and Usage
SG24-7676-00
Draft Redbooks, last update 8 January 2009
http://www.redbooks.ibm.com/redbooks/pdfs/sg245250.pdf

IBM System Storage Solutions Handbook, SG24-5250-06
Redbooks, published 12 October 2006, last updated 29 May 2007
http://www.redbooks.ibm.com/redbooks/pdfs/sg245250.pdf

iSeries and IBM TotalStorage: A Guide to Implementing External Disk on eServer i5, SG24-7120-00, *Redbooks*, published 16 October 2005
http://www.redbooks.ibm.com/redbooks/pdfs/sg247120.pdf