Performance Evaluation of DR550 v4.5 Using the File System Gateway (FSG)

March 2, 2009

Author Mike Roll, Vernon Miller
Manager Kamran Amini
Modular Storage and Power Blade Performance, Tucson
IBM SYSTEMS AND TECHNOLOGY GROUP

NOTICES AND DISCLAIMERS

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice.

The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe on IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

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

IBM, System Storage, and AIX are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

NOTE:

The hard copy version of this document is FOR REFERENCE ONLY.

It is the responsibility of the user to ensure that they have the current version. Any outdated hard copy is invalid and must be removed from possible use. It is also the responsibility of the user to ensure the completeness of this document prior to use.

Table of Contents

DR550 v4.5 File System Gateway (FSG) Performance Evaluation

Introduction

The IBM System Storage DR550, one of the IBM Data Retention offerings, is an integrated offering for clients that need to retain and preserve electronic business records. It is designed to help store, retrieve, manage and retain regulated and non-regulated data. In addition to managing compliance data, it can also be an archiving solution for other types of data.

However, the main focus of the DR550 is to help companies manage (archive and retrieve) data assets in a non-erasable and non-rewritable storage system. It provides a secure data repository where deletion or modification of data is completely disallowed except through a well defined retention and expiration policy. New industry and government regulations are leading to greater amounts and types of data that need to be stored in this type of format.

DR550 Product Description

The DR550 brings together off-the-shelf IBM hardware and software products. The hardware comes already mounted in a secure rack, including IBM's POWER5+ servers with AIX, IBM System Storage Archive Manager (SSAM), and the IBM System Storage DS4700 Express1.

The DR550 is fully integrated. The hardware has been installed, cabled and tested in a rack, and the software has been installed and largely pre-configured. System administrators are required to define the SSAM management classes associated with the retention periods for their business needs. When combined with a content management application, data can be managed throughout it's lifecycle, including automatic deletion at the end of it's retention period.

Purpose of This Performance Evaluation

The purpose of this Performance Evaluation was to measure the performance of archiving and retrieving objects to and from the DR550 across the File System Gateway (FSG). The size of the objects were varied (512 byte, 1K, 2K, 4K & 8K) along with the number of concurrent application threads (1, 2, 4, 8, 16 & 32) that either archived or retrieved them.

File System Gateway (FSG)

The DR550 v4.5 supports an optional File System Gateway server that enables client applications to perform file-based stores and retrieves from the DR550 through either the Common Internet File System (CIFS) or Network File System (NFS) protocols. The File System Gateway interacts with the DR550 over the System Storage Archive Manager (SSAM) API.

As files are ingested through the File System Gateway, they are cached locally and forwarded to the DR550 for persistent storage. The DR550 File System Gateway (DRG) service is the software module of the File System Gateway that streams data to the DR550. The DRG service also caches a copy of the data to speed its retrieval.

When a client application requests an object that must be retrieved from the DR550, the File System Gateway requests the object from the DR550. If the object is cached in the storage subsystem, then it's sent directly to the File System Gateway. Otherwise, the File System Gateway requests the object from the Archive Service (ARC) which is middleware that retrieves the object from the DR550 and sends it to the File System Gateway. The File System Gateway then verifies the object, caches it in the storage subsystem, and forwards it.

Performance Workload Tools

The Flexible File System Benchmark (FFSB v5.2.1) is an open source network-based application workload generation tool that runs in a Linux environment from an NFS mount point to generate a variety of workloads across the DR550 File System Gateway (FSG). For each run, it captures the performance information and generates a performance report.

Performance Measurement Approach

The FFSB tool was used to generate files that were either stored or retrieved through the File System Gateway. Each performance run used 100 40MB files and lasted about five minutes. The throughput for each run was captured in both MB per second and IO operations per second against the following variables:

- Increasing the file transfer size of each IO (512 Bytes, 1K, 2K, 4K, & 8K)
- Varying the number of concurrent application threads from 1, 2, 4, 8, 16 & 32
- Running with the FSG cache on & off

FFSB was run on Linux client against the File System Gateway NFS mount point which was specified as a directory. Random read & write behaviors were implemented by the FFSB tool by either reading or writing individual blocks starting from a random point on the file, continuing until the entire amount specified had been processed. Sequential read/write behavior meant that blocks were either read or written starting from the first byte in the file.

Performance Measurement Environment

The System Storage Archive Manager was embedded on an IBM System P5 52A using POWER5+ processors. This entry-level server incorporates a number of the attributes of IBM's high-end servers.

The File System Gateway ran on the 2229-FSG File System Gateway node with 2 GB of RAM, powered by a 3.2GHz dual core Xeon processor using SUSE Linux™ Enterprise Server Version 10.
The 2229-FSG connected to the customer network using either 1 Gbps Ethernet adapters. The 2229-FSG connected to the DR550 using a 10/100/1000 TX Ethernet adapter. As seen below, one bonded Ethernet connection was used between the customer network and the DR550's SSAM server.

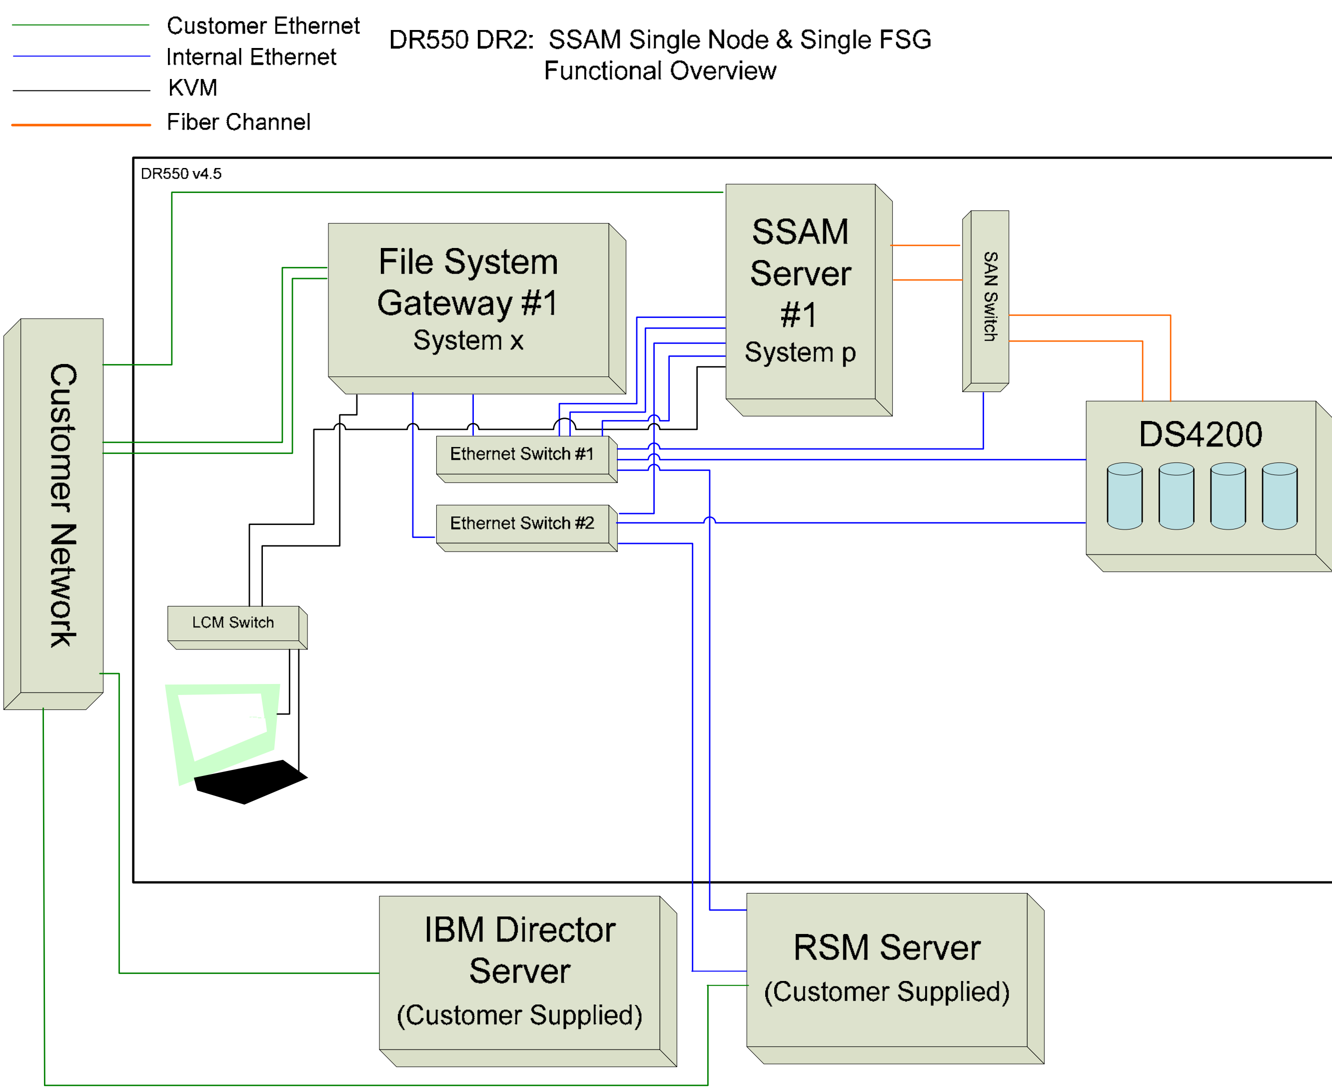

Figure #1 – Logical Connection Diagram of the DR550 DR2 used for the performance evaluation.

The storage subsystem used in the DR550 was the DS4200 with 16 750 GB SATA drives in the control unit and two EXP420 16-disk expansion drawers. There were 48 disk drives total configured with over 26 TB of storage. The performance evaluation was performed with both write cache and write cache mirroring enabled for the logical drives. The DS4200 was attached to the SSAM server through a switch using two 4Gb/second Fiber Channel lines.

The primary software used to drive the DR550 is IBM's System Storage Archive Manager (SSAM). It's designed to help customers protect the integrity of data as well as automatically enforce data retention policies. For this performance evaluation, archived data was written to a single SSAM storage pool which was spread across all the disk arrays on the DS4200.

Using policy-based management, data can be stored indefinitely, can be expired based on a retention event, or have a predetermined expiration date. In addition, the retention enforcement feature may be applied to data using deletion hold and release interfaces which hold data for an indefinite period of time, regardless of the expiration date or defined event.

The policy software is also designed to prevent modifications or deletions after the data is stored. With support for open standards, the new technology is designed to provide customers flexibility to use a variety of content management or archive applications.

1. Random Reads/Retrieves without Cache (MB/sec)

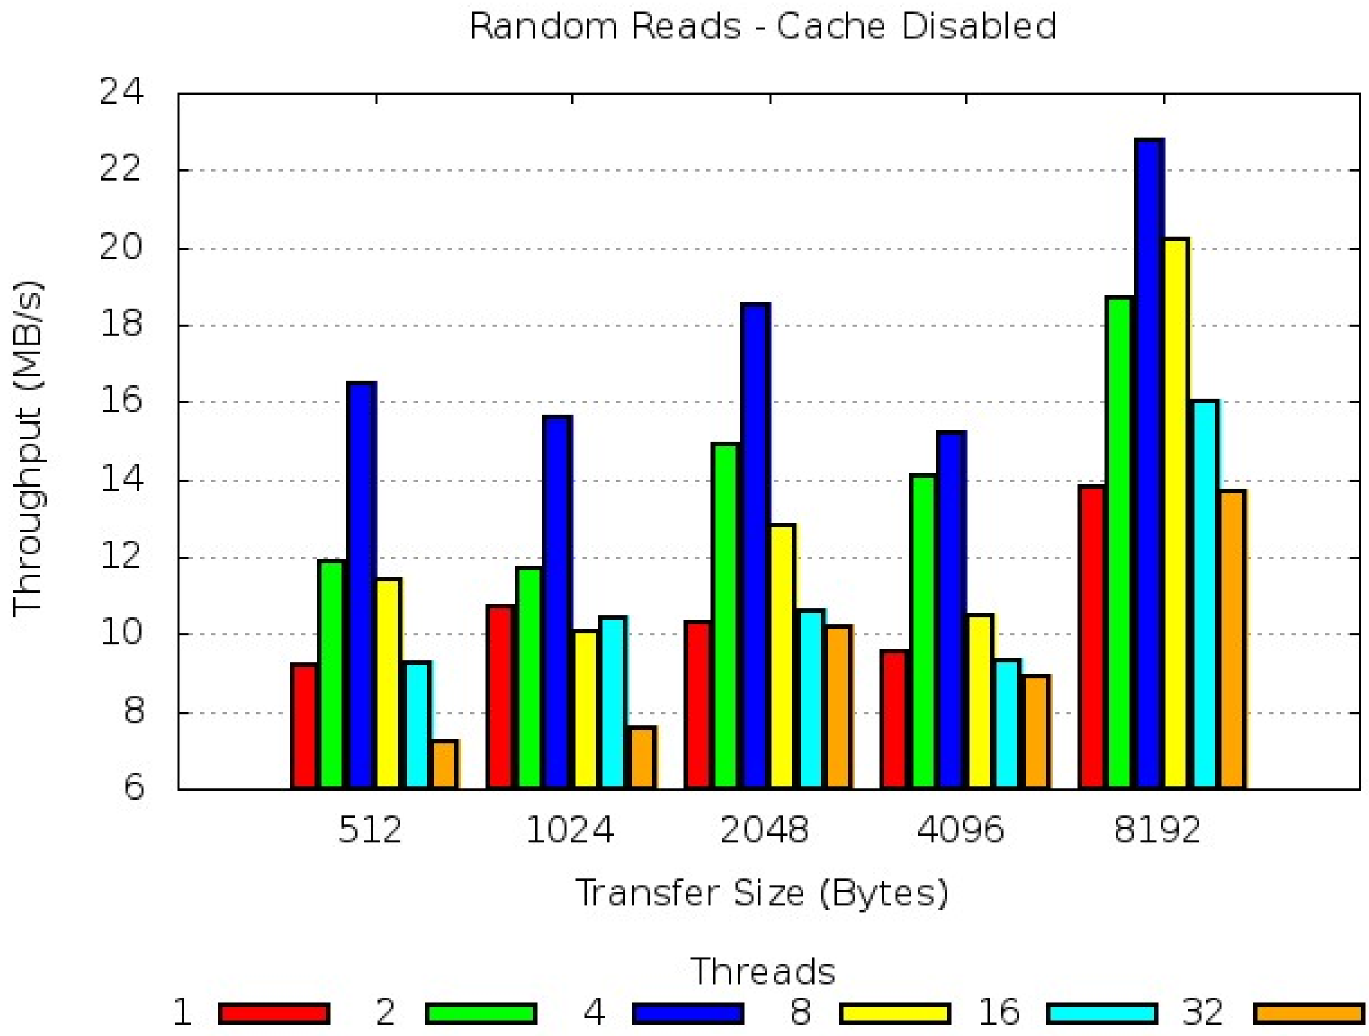

This diagram describes how the random RETRIEVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads also successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow reads directly against the DR550 without any caching in the File System Gateway.

Four concurrent application threads achieved the highest throughputs for each of the different transfer sizes, with 23 MB/sec as the highest throughput .

2. Random Reads/Retrieves without Cache (IOPS)

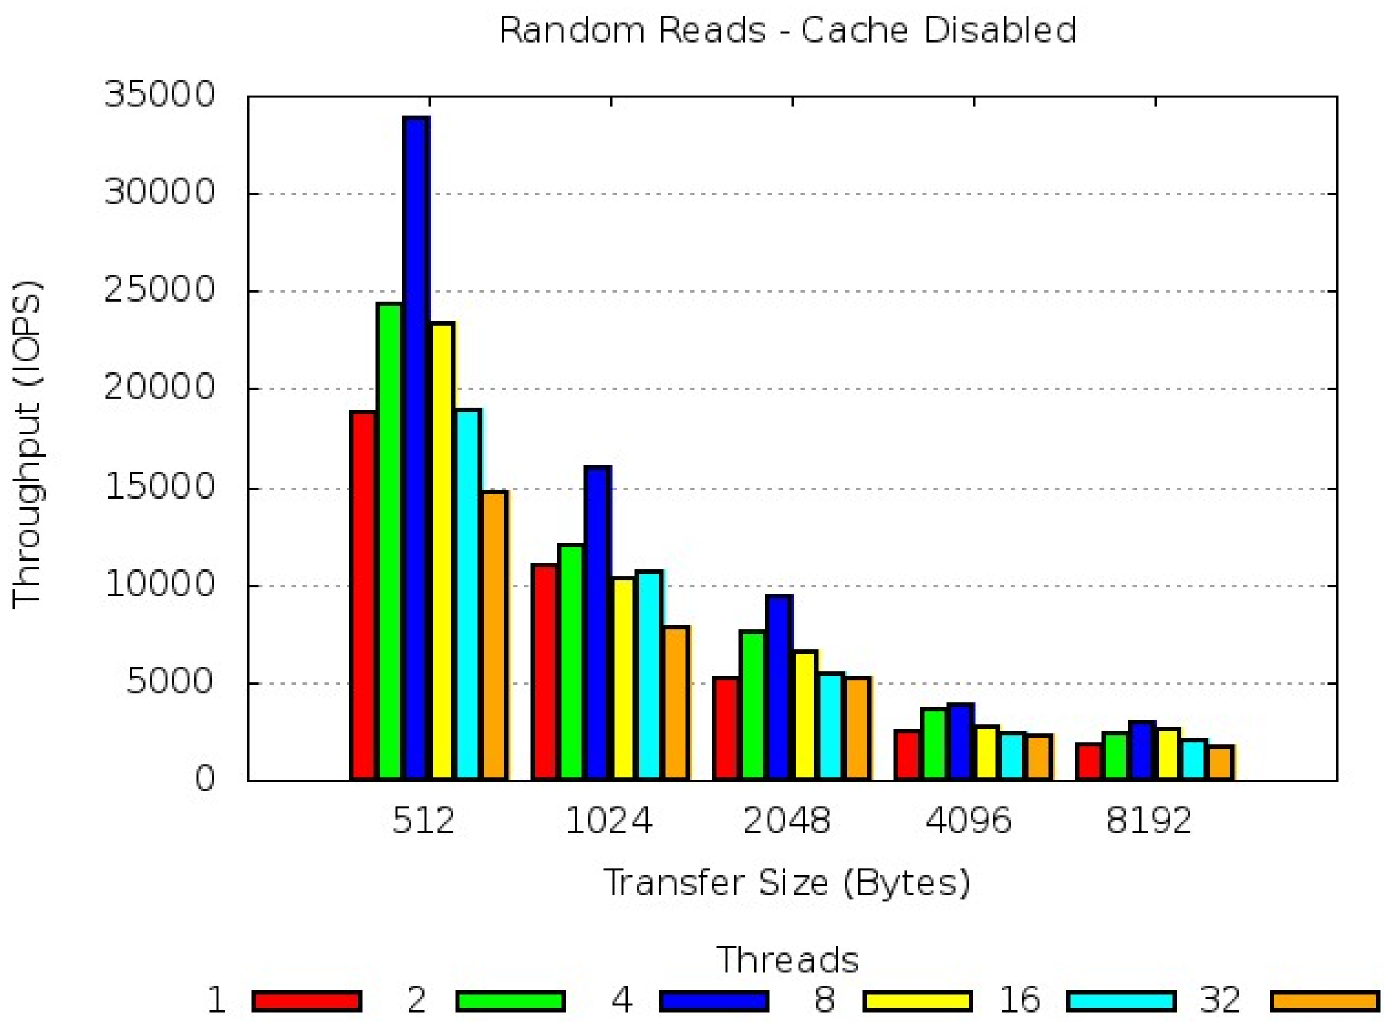

This diagram describes how the random RETRIEVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads were successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow reads directly against the DR550 without any caching in the File System Gateway.

Four concurrent application threads achieved the highest throughputs for each of the different transfer sizes with the largest file transfer size showing the fewest IO's/sec.

3. Random Reads/Retrieves with Cache (MB/sec)

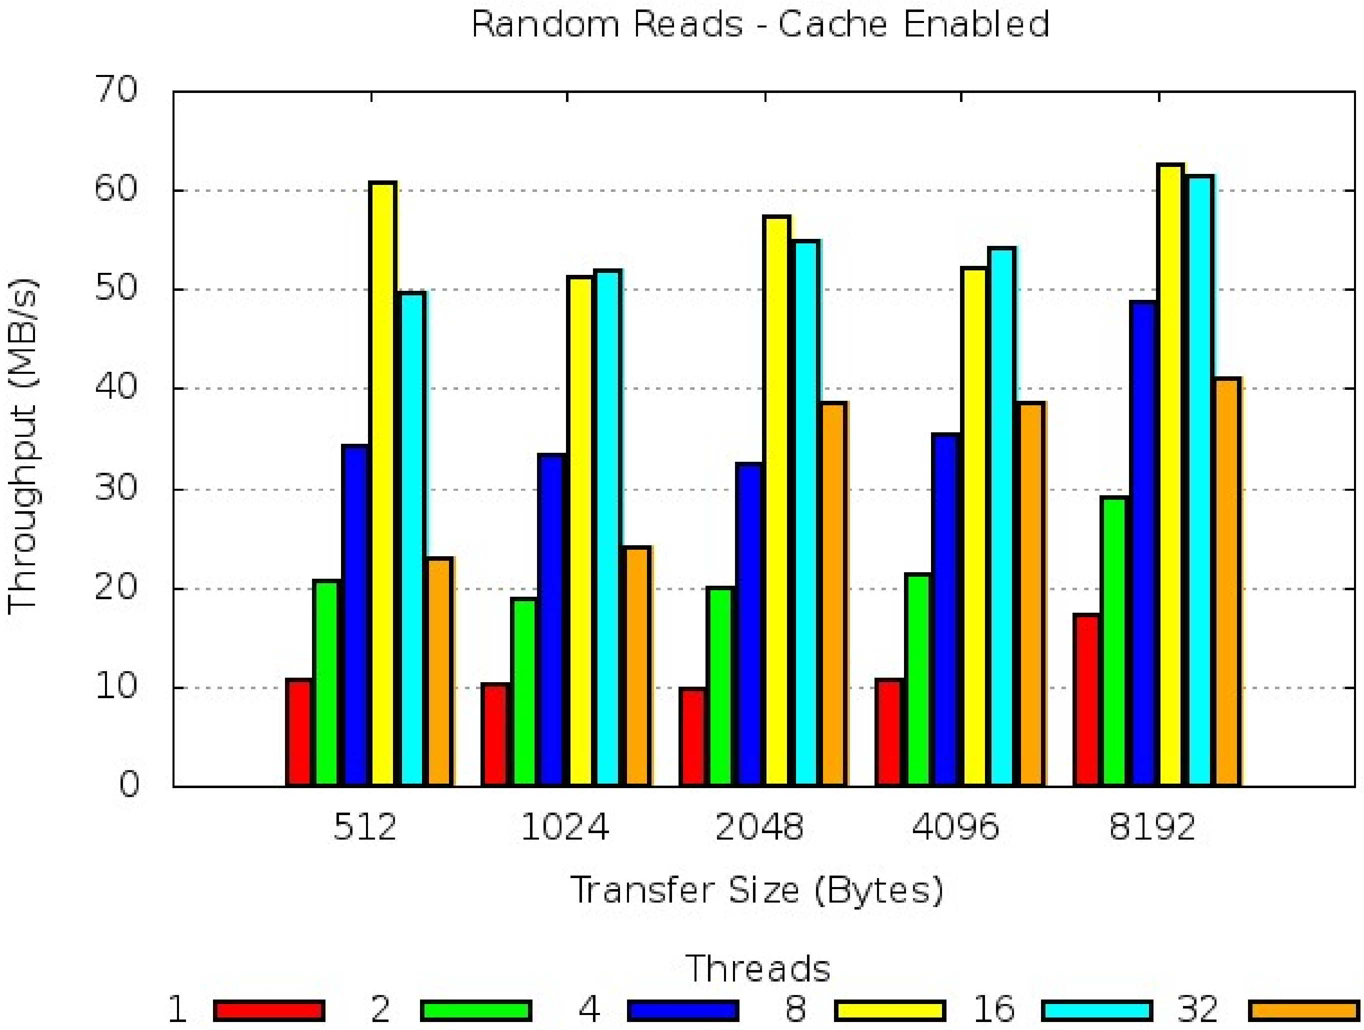

This diagram shows how the random RETRIEVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads were also successively increased.

The File System Gateway cache was enabled, allowing reads to leverage the cache in the File System Gateway to achieve greater throughput when retrieving data from the DR550.

With FSG Cache being used, a higher number of concurrent application threads (either 8 or 16) are required to achieve the maximum throughput of 63 MB/sec.

4. Random Reads/Retrieves with Cache (IOPS)

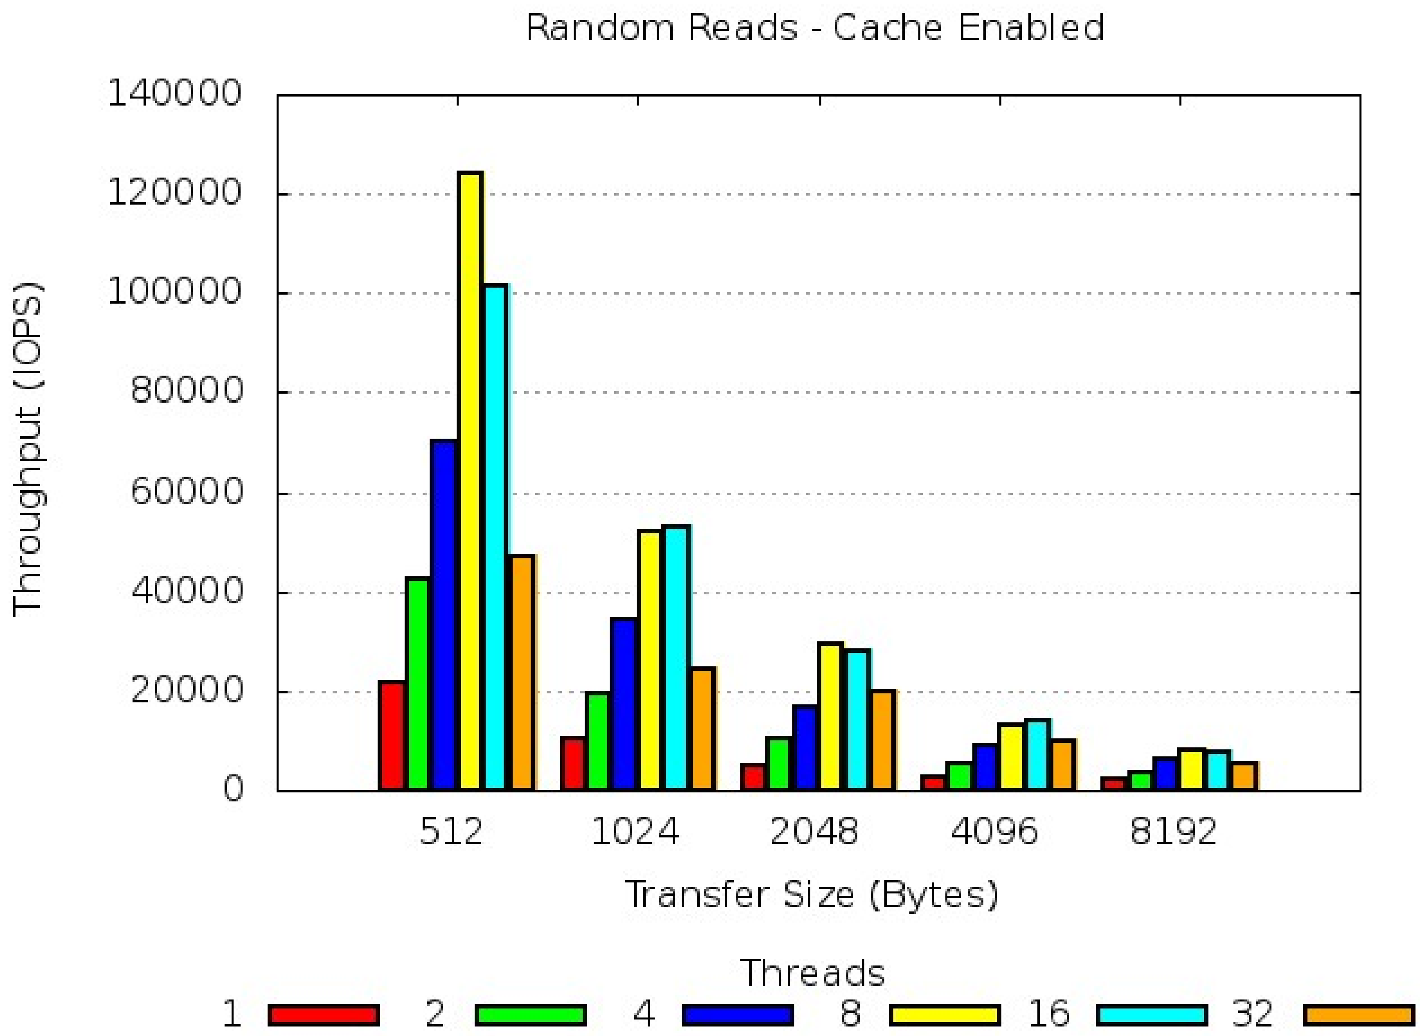

This diagram again describes how the random RETRIEVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads were successively increased.

The File System Gateway cache was enabled, allowing reads to leverage the cache in the File System Gateway to achieve greater throughput when retrieving data from the DR550.

Again, we see that the smallest file transfer size (512 bytes) providing the highest number of IO's/sec.

5. Random Writes/Archives without Cache (MB/sec)

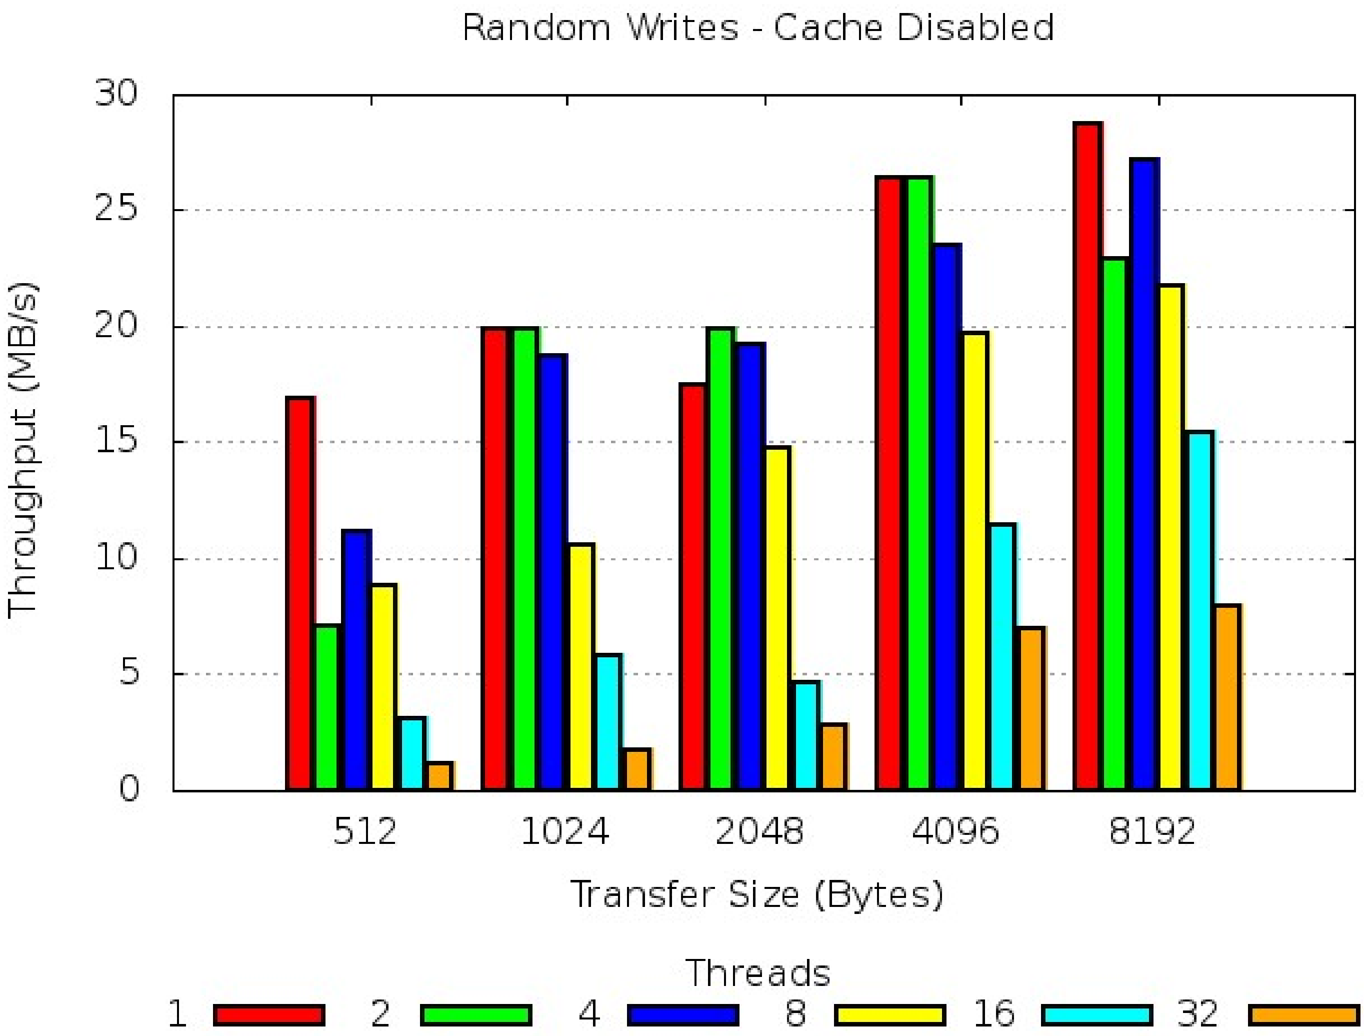

This diagram describes how the ARCHIVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads also successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow writes directly against the DR550 without any caching in the File System Gateway.

For random writes without FSG cache, we see that a single application thread generally provides the highest throughput for each data transfer size, with a maximum throughput of 28 MB/sec.

6. Random Writes/Archives without Cache (IOPS)

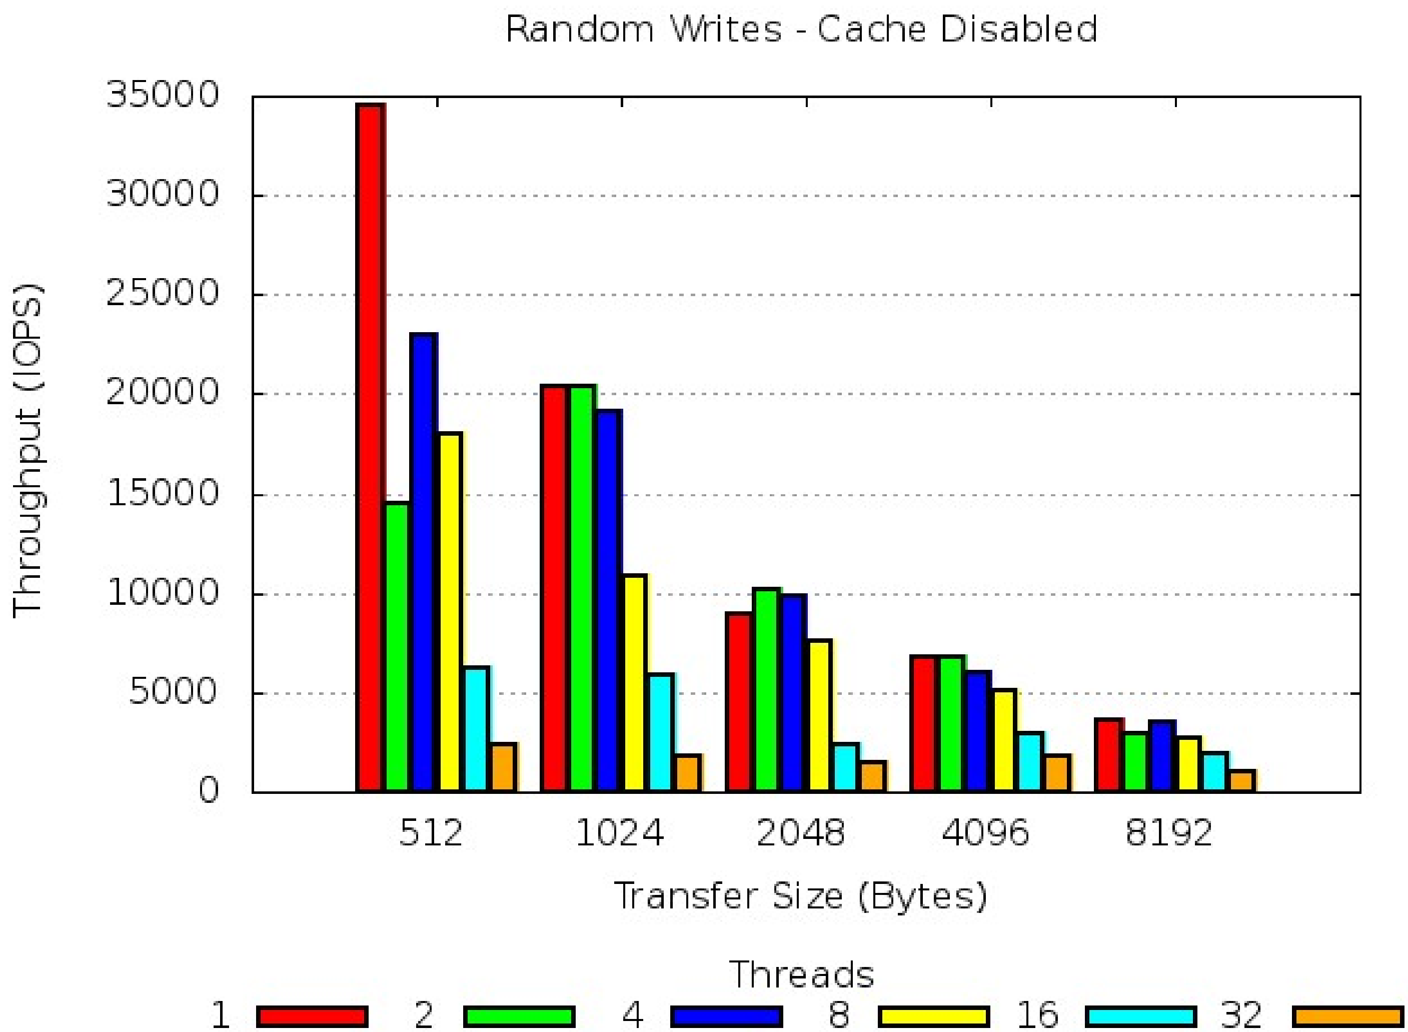

This diagram describes how the ARCHIVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads also successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow writes directly against the DR550 without any caching in the File System Gateway.

For the smallest file transfer size (512 bytes), the single application thread has an outsize impact on the the highest number of IO's/sec, with fewer IO's as the file transfer size increases.

7. Random Writes/Archives with Cache (MB/sec)

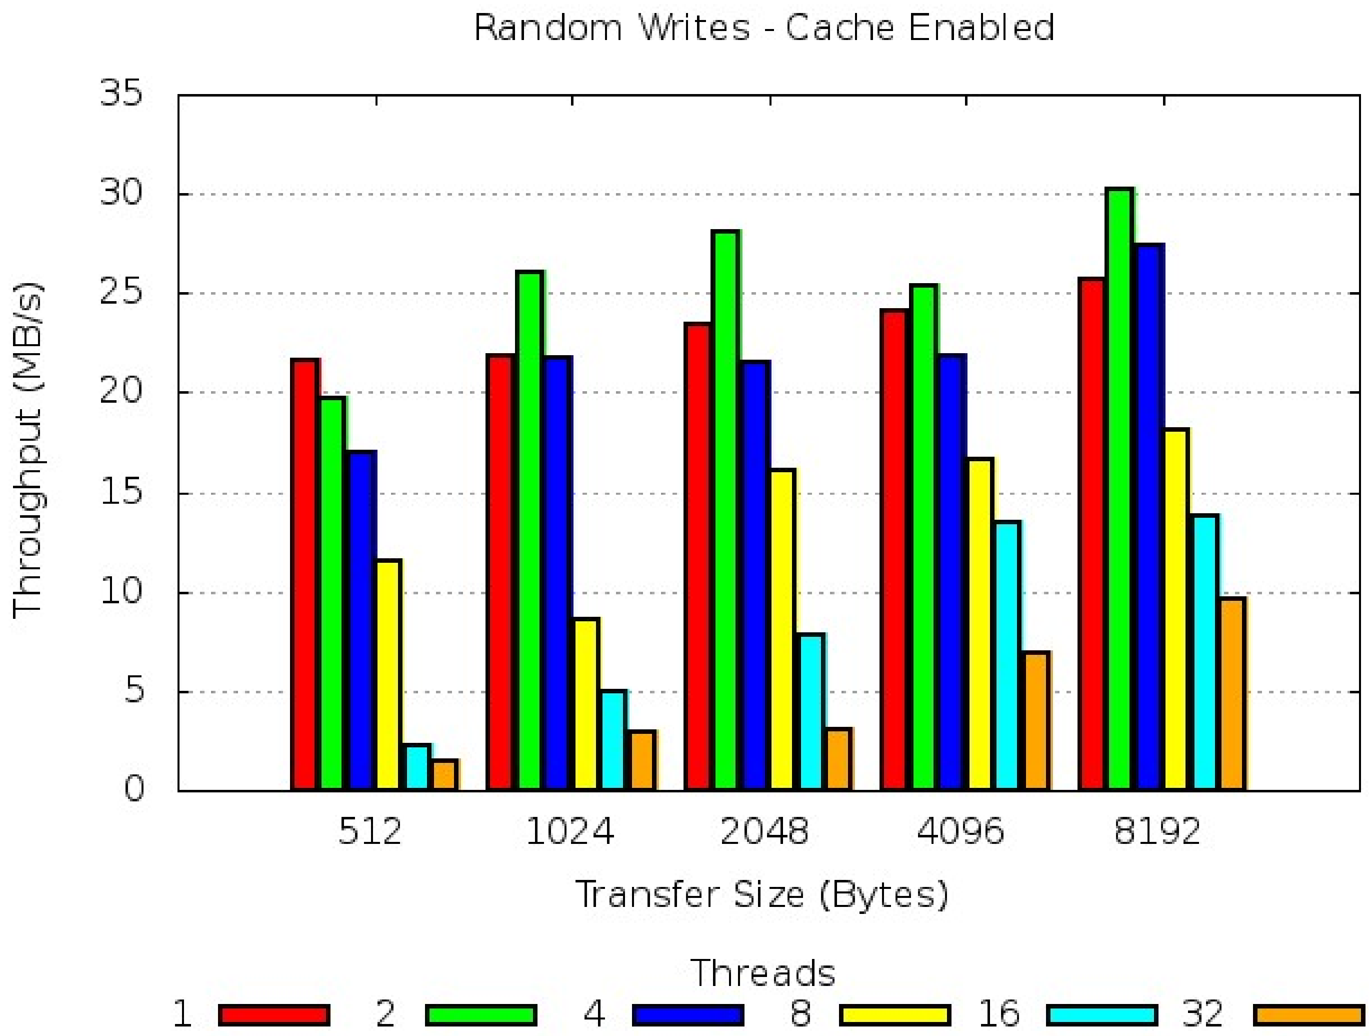

This diagram shows how the ARCHIVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads were also successively increased.

The File System Gateway cache was enabled, allowing writes to leverage the cache in the File System Gateway to achieve greater throughput when archiving data to the DR550.

Random writes using FSG cache require two concurrent application threads to achieve the highest throughput (31 MB/sec).

8. Random Writes/Archives with Cache (IOPS)

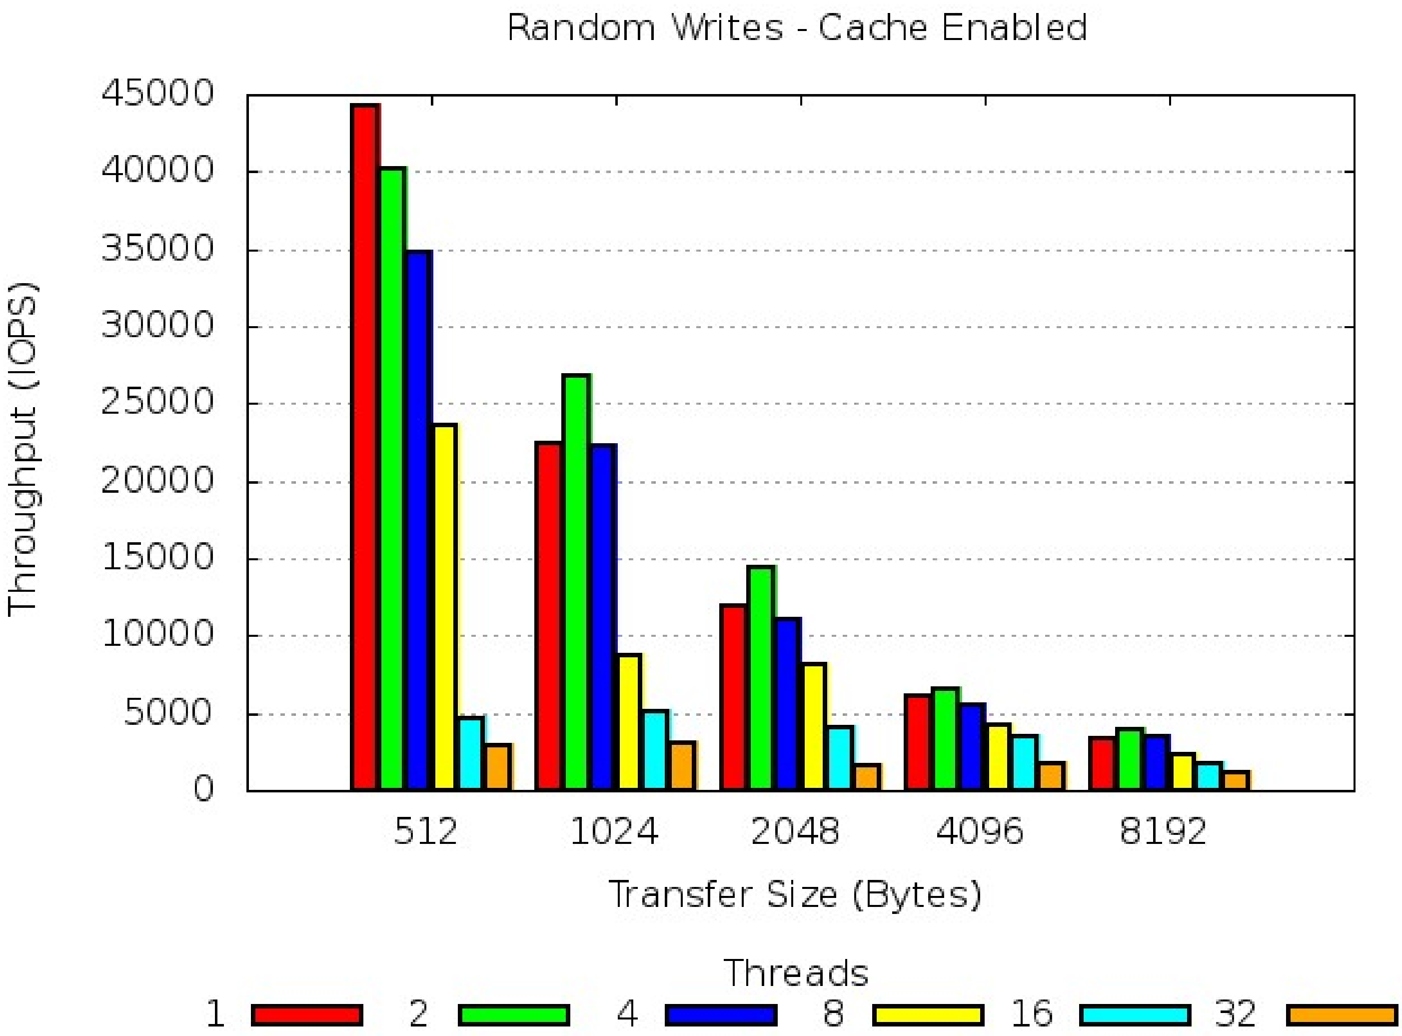

This diagram shows how the ARCHIVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads were also successively increased.

The File System Gateway cache was enabled, allowing writes to leverage the cache in the File System Gateway to achieve greater throughput when archiving data to the DR550.

For the smallest file transfer size (512 bytes), the single application thread shows the highest number of IO's/sec, followed by 2 application threads and fewer IO's/sec as the file transfer size increases.

9. Sequential Reads without Cache (MB/sec)

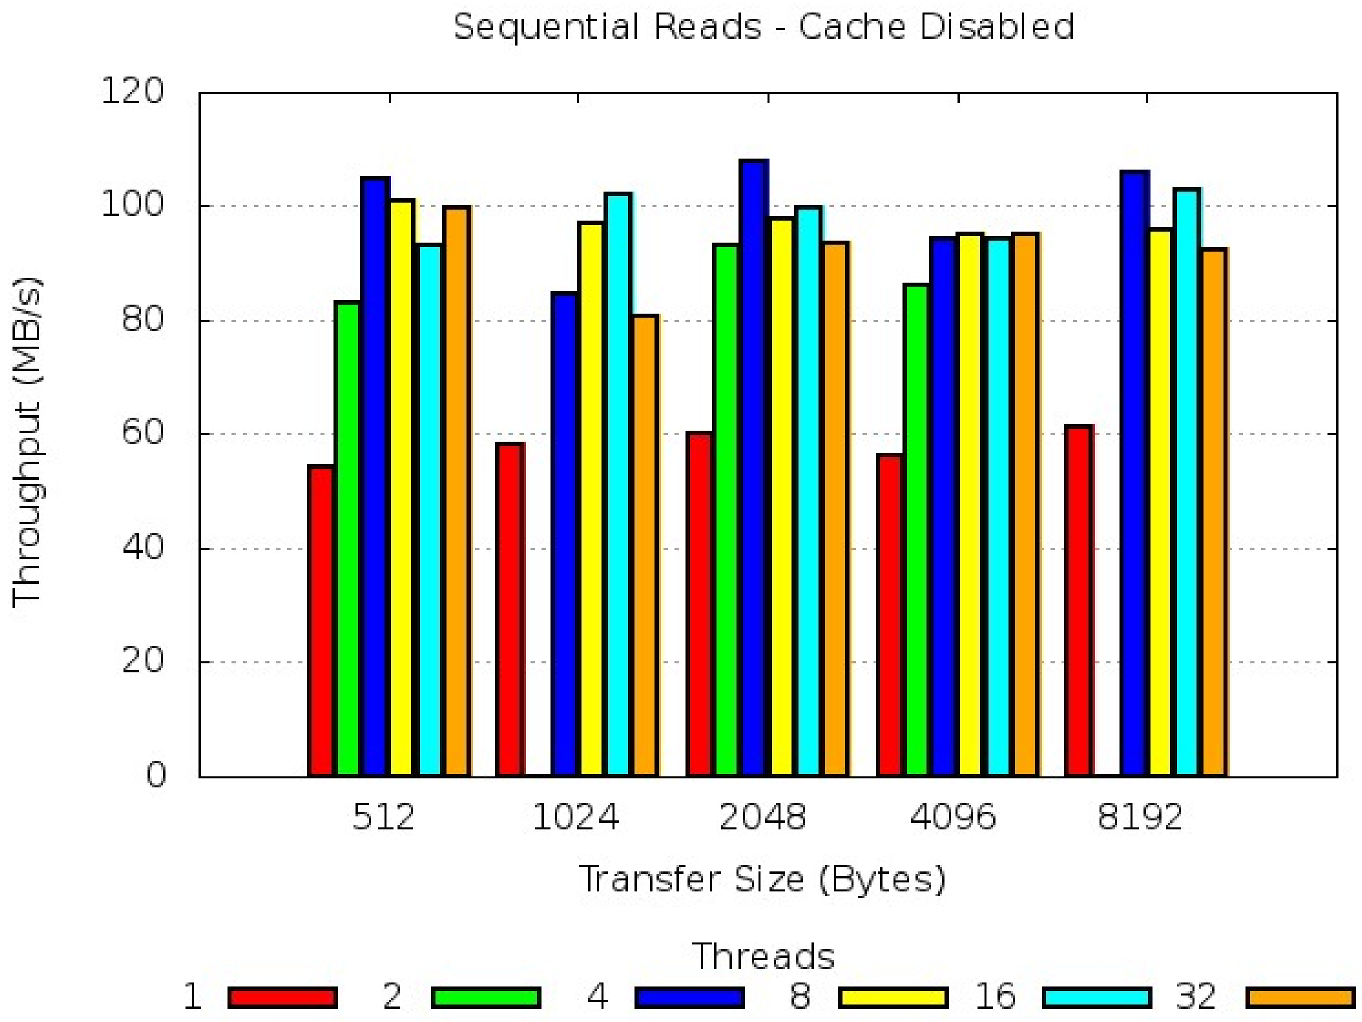

This diagram describes how a sustained, sequential RETRIEVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads also successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow reads directly against the DR550 without any caching in the File System Gateway.

Four concurrent application threads achieved the highest throughputs for each of the different transfer sizes, with 110 MB/sec as the highest throughput .

10. Sequential Reads without Cache (IOPS)

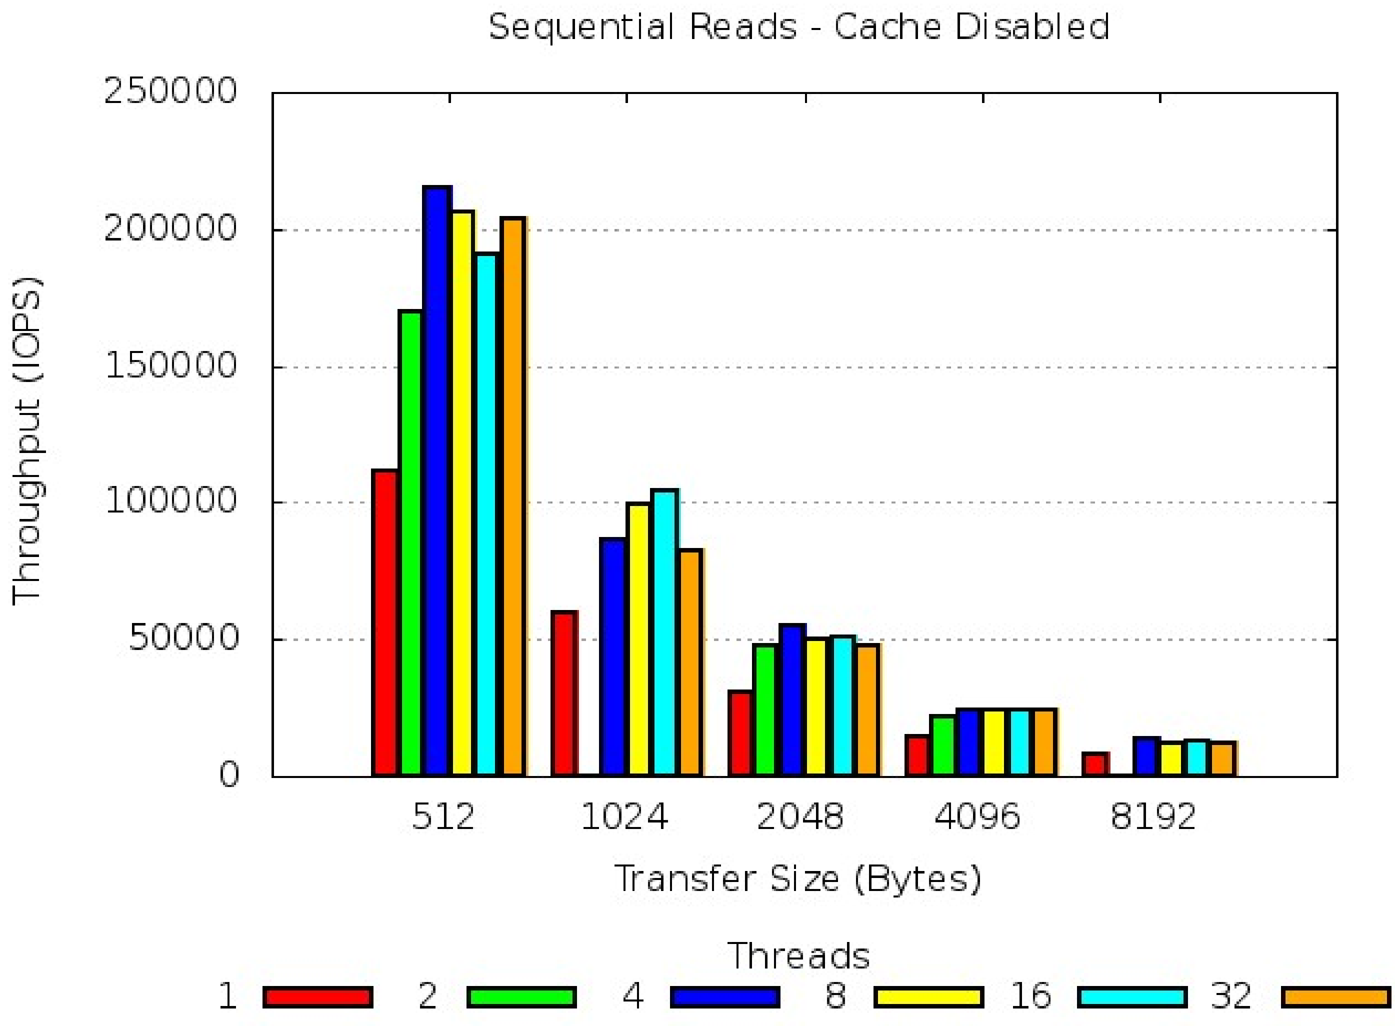

The diagram above shows how a sustained, sequential RETRIEVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads were successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow reads directly against the DR550 without any caching in the File System Gateway.

Four concurrent application threads achieved the highest throughputs for each of the different transfer sizes with the largest file transfer size showing the fewest IO's/sec.

11. Sequential Reads with Cache (MB/sec)

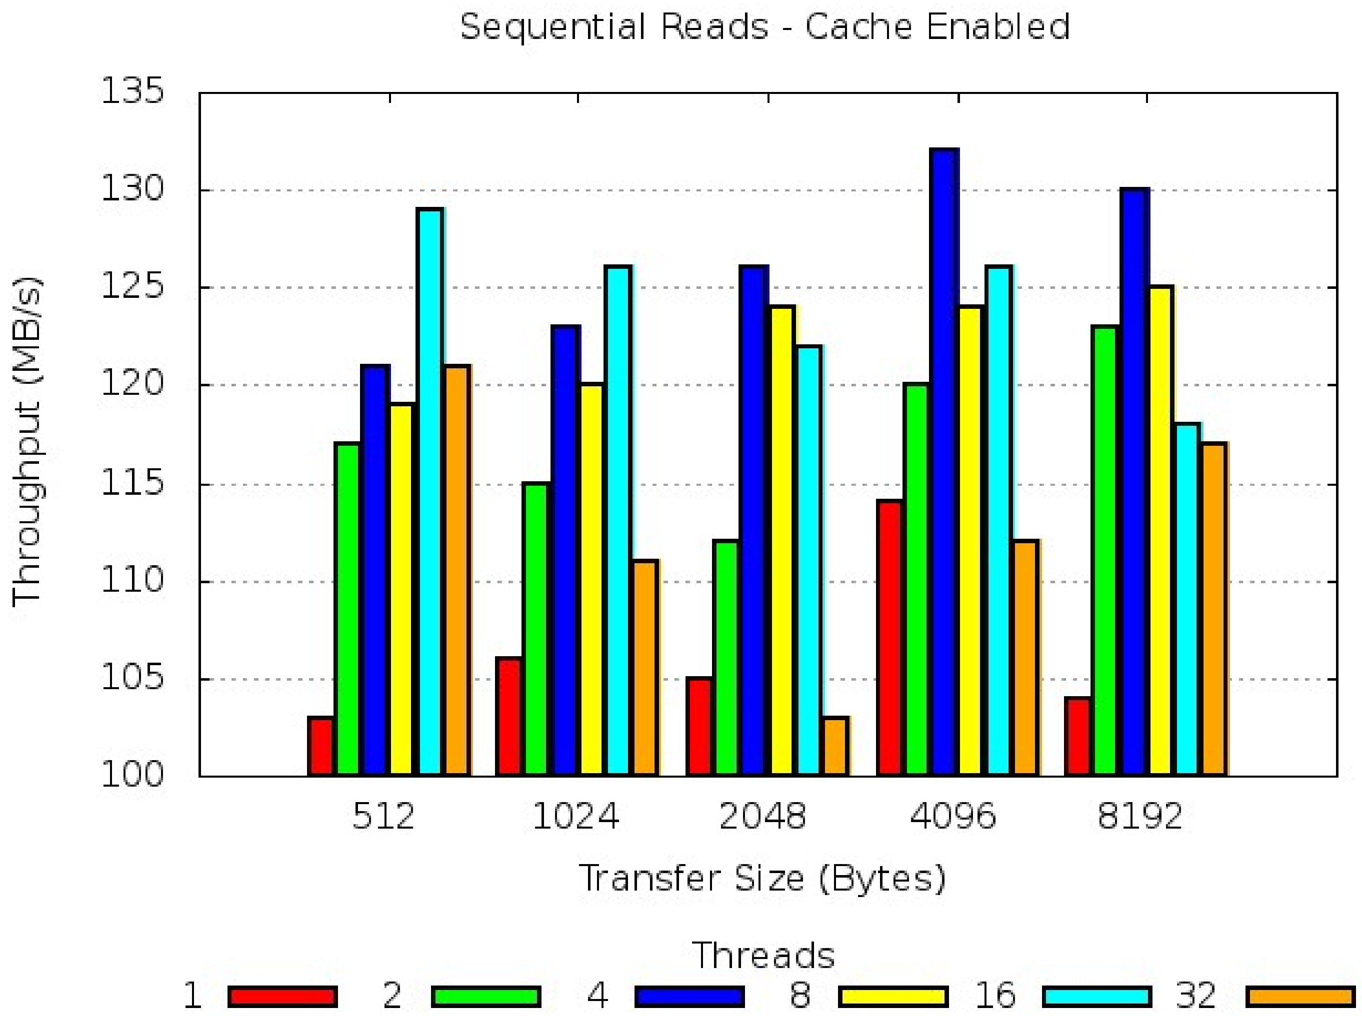

This diagram illustrates how a sequential, sustained RETRIEVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads were also successively increased.

The File System Gateway cache was enabled, allowing reads to leverage the cache in the File System Gateway to achieve greater throughput when retrieving data from the DR550.

With FSG Cache being used, a higher number of concurrent application threads are required to achieve maximum throughput.

12. Sequential Reads with Cache (IOPS)

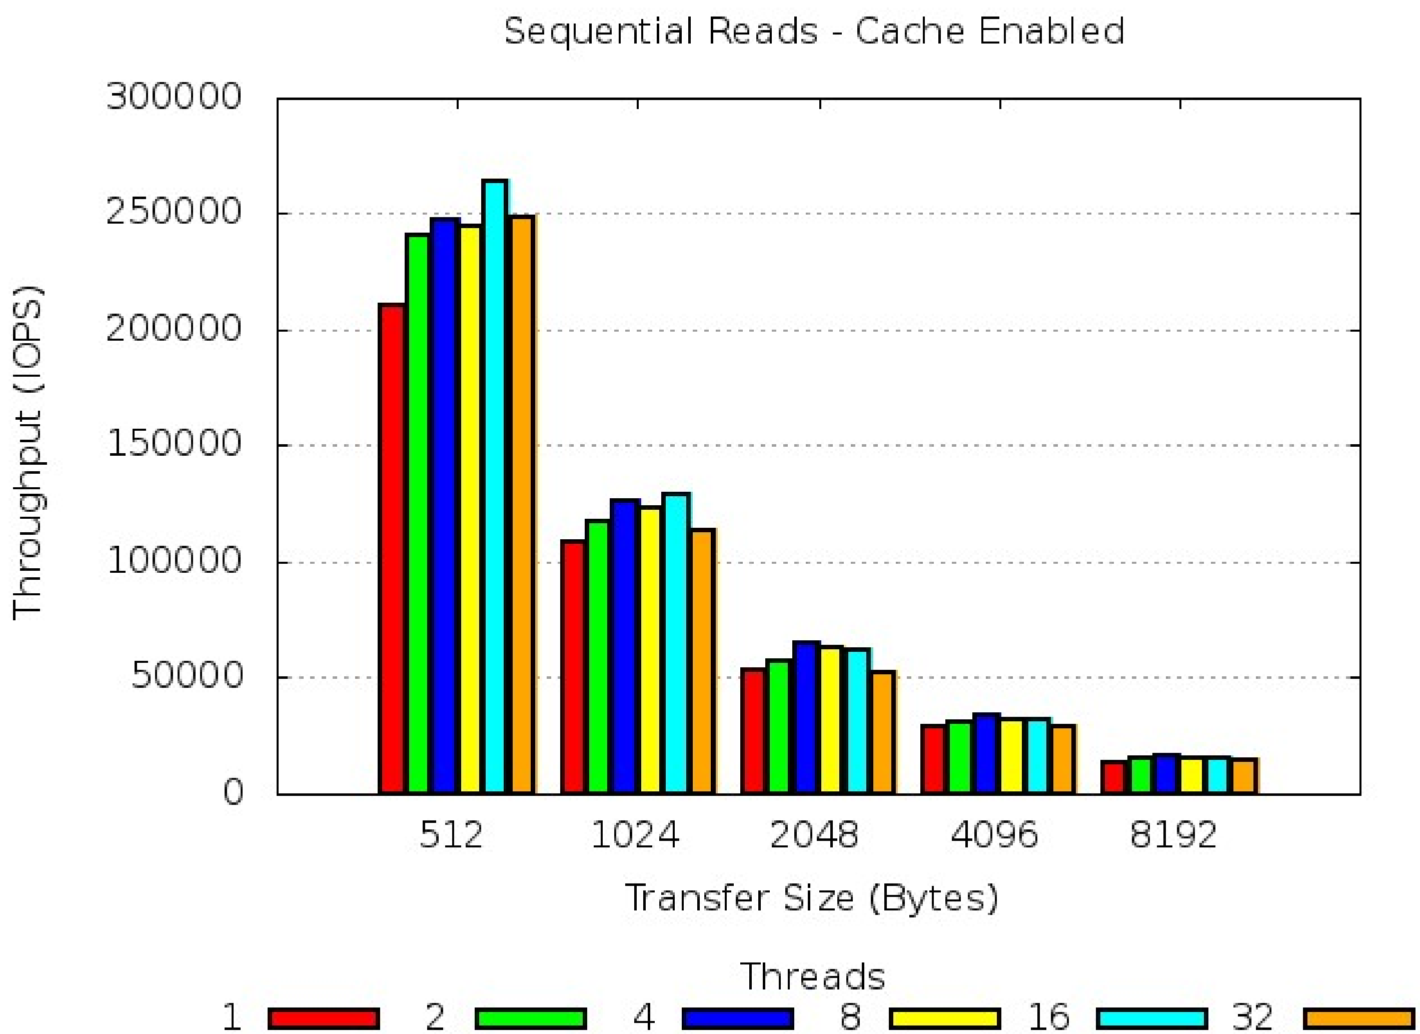

This diagram again describes how a sustained, sequential RETRIEVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being retrieved from the DR550 increased and the number of RETRIEVE application threads were successively increased.

The File System Gateway cache was enabled, allowing reads to leverage the cache in the File System Gateway to achieve greater throughput when retrieving data from the DR550.

Again, we see that the smallest file transfer size (512 bytes) providing the highest number of IO's/sec.

13. Sequential Writes without Cache (MB/sec)

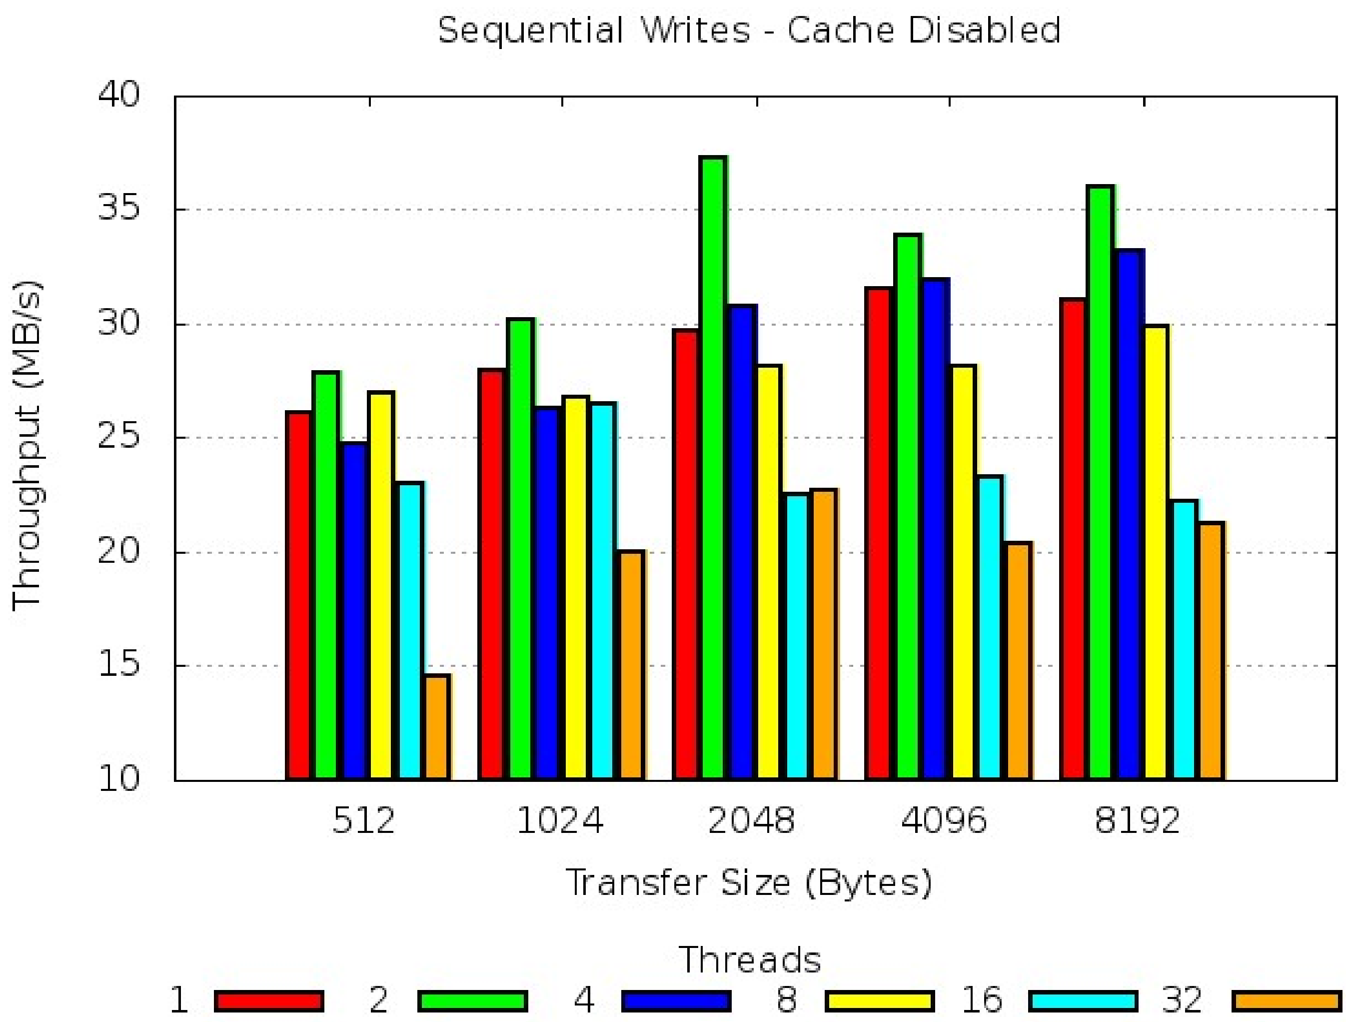

This diagram shows how a sustained sequential, ARCHIVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads also successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow writes directly against the DR550 without any caching in the File System Gateway.

For sequential writes without FSG cache, we see that two concurrent application threads provide the highest throughput for each data transfer size, with a maximum throughput of 37 MB/sec.

14. Sequential Writes without Cache (IOPS)

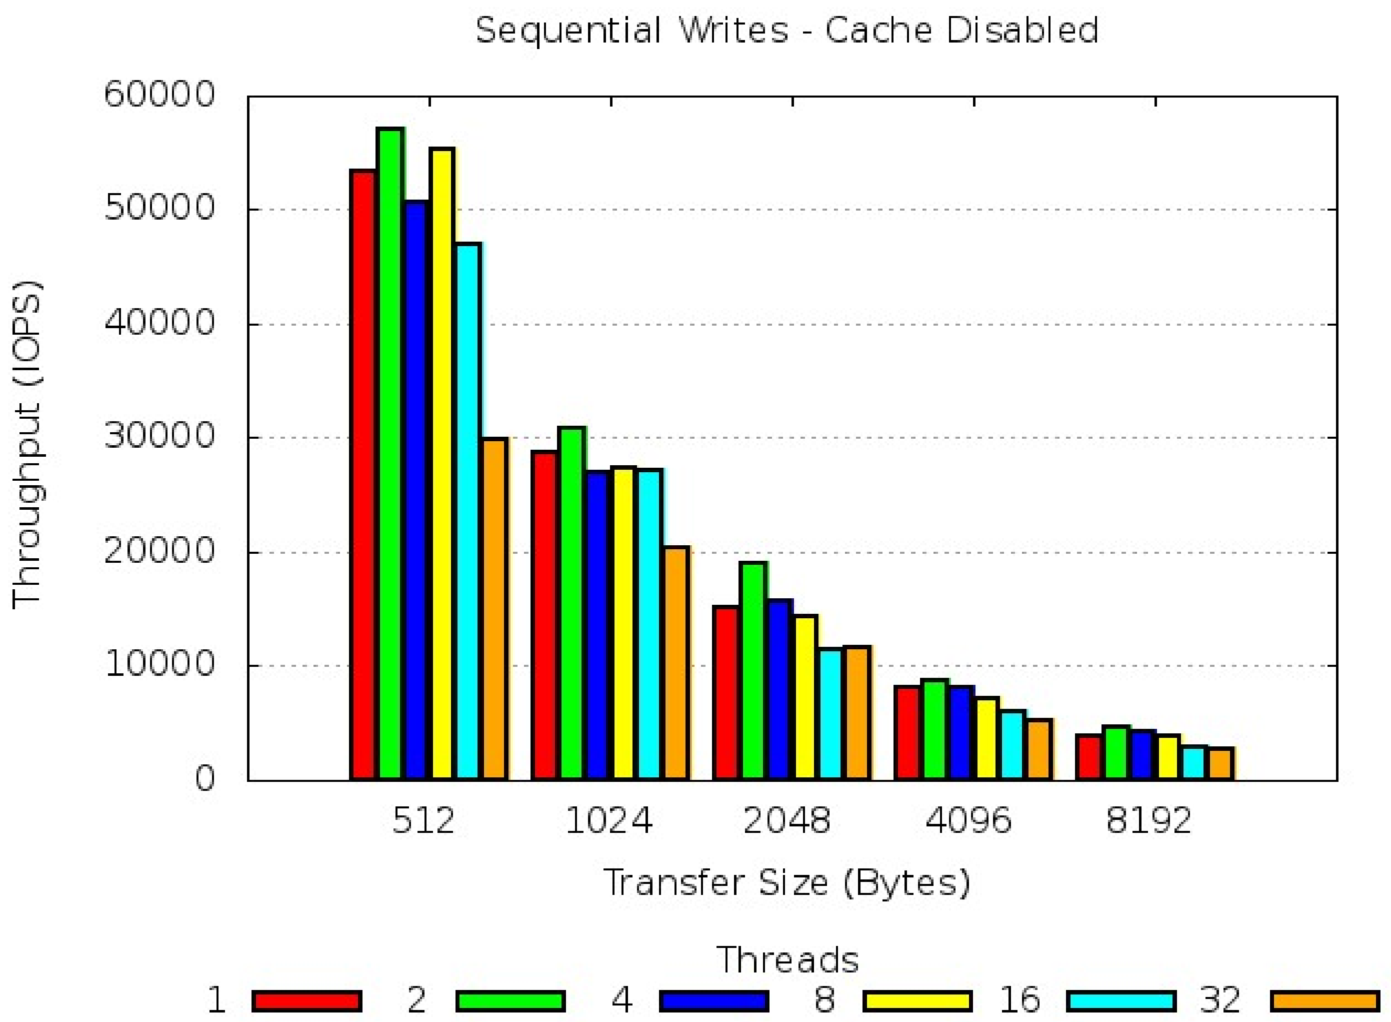

This diagram describes how the sustained, sequential ARCHIVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads also successively increased.

For these measurements, the cache on the File System Gateway was disabled to allow writes directly against the DR550 without any caching in the File System Gateway.

Once again, two concurrent application threads provide the highest number of IO's/sec for each data transfer size, with the throughput decreasing as the transfer size increases.

15. Sequential Writes with Cache (MB/sec)

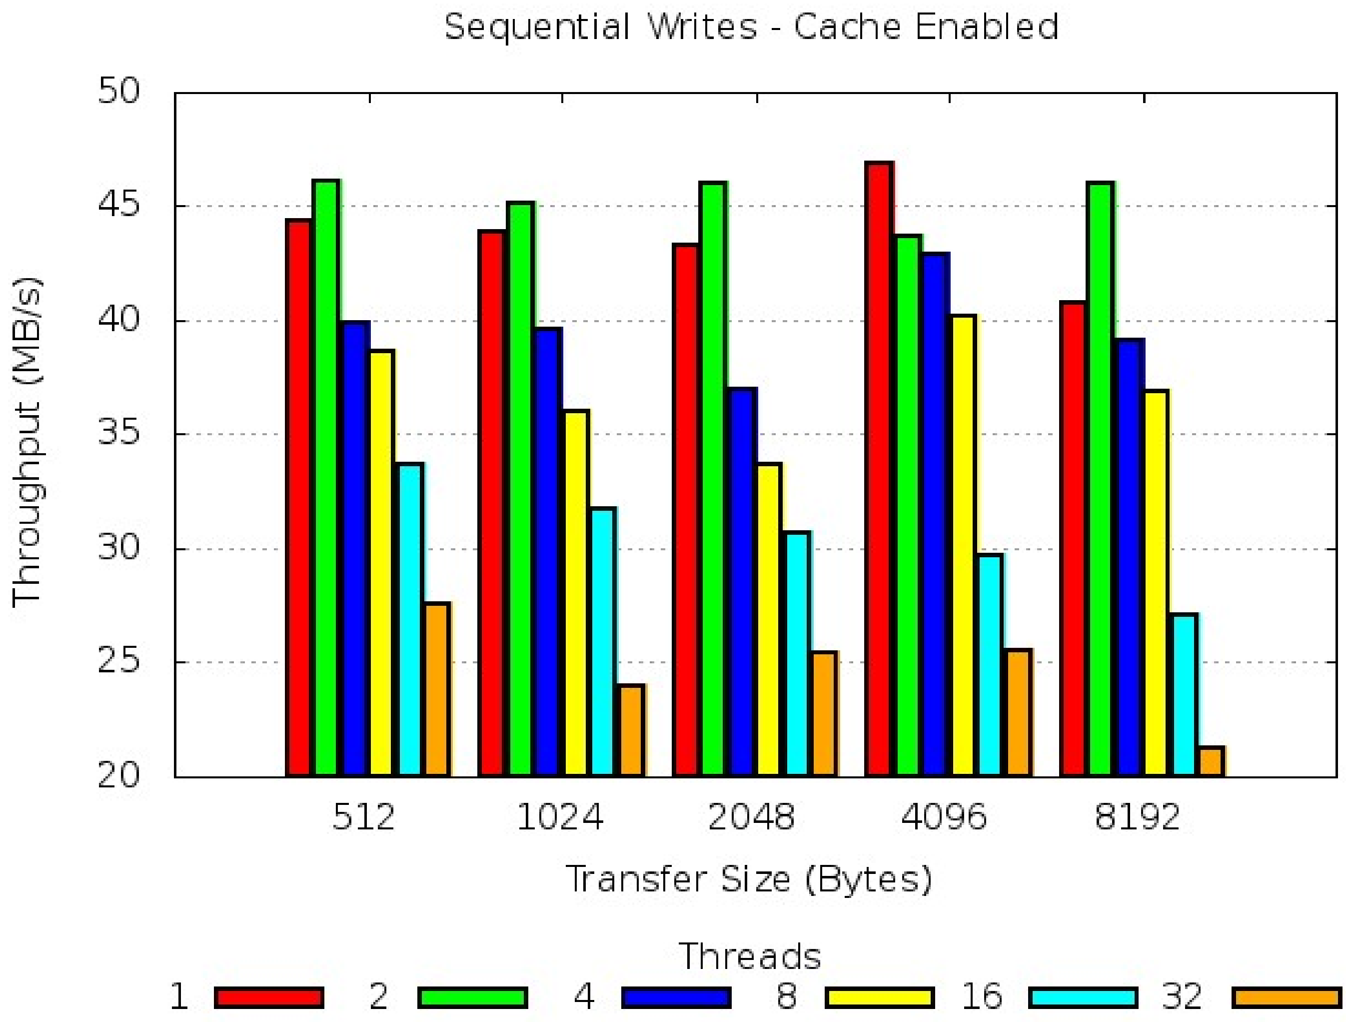

The diagram above shows how the sustained, sequential ARCHIVE data throughput rate (MB/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads were also successively increased.

The File System Gateway cache was enabled, allowing writes to leverage the cache in the File System Gateway to achieve greater throughput when archiving data to the DR550.

Sequential writes using FSG cache generally require two concurrent application threads to achieve higher throughput, although one thread generated the highest at 47 MB/sec.

16. Sequential Writes with Cache (IOPS)

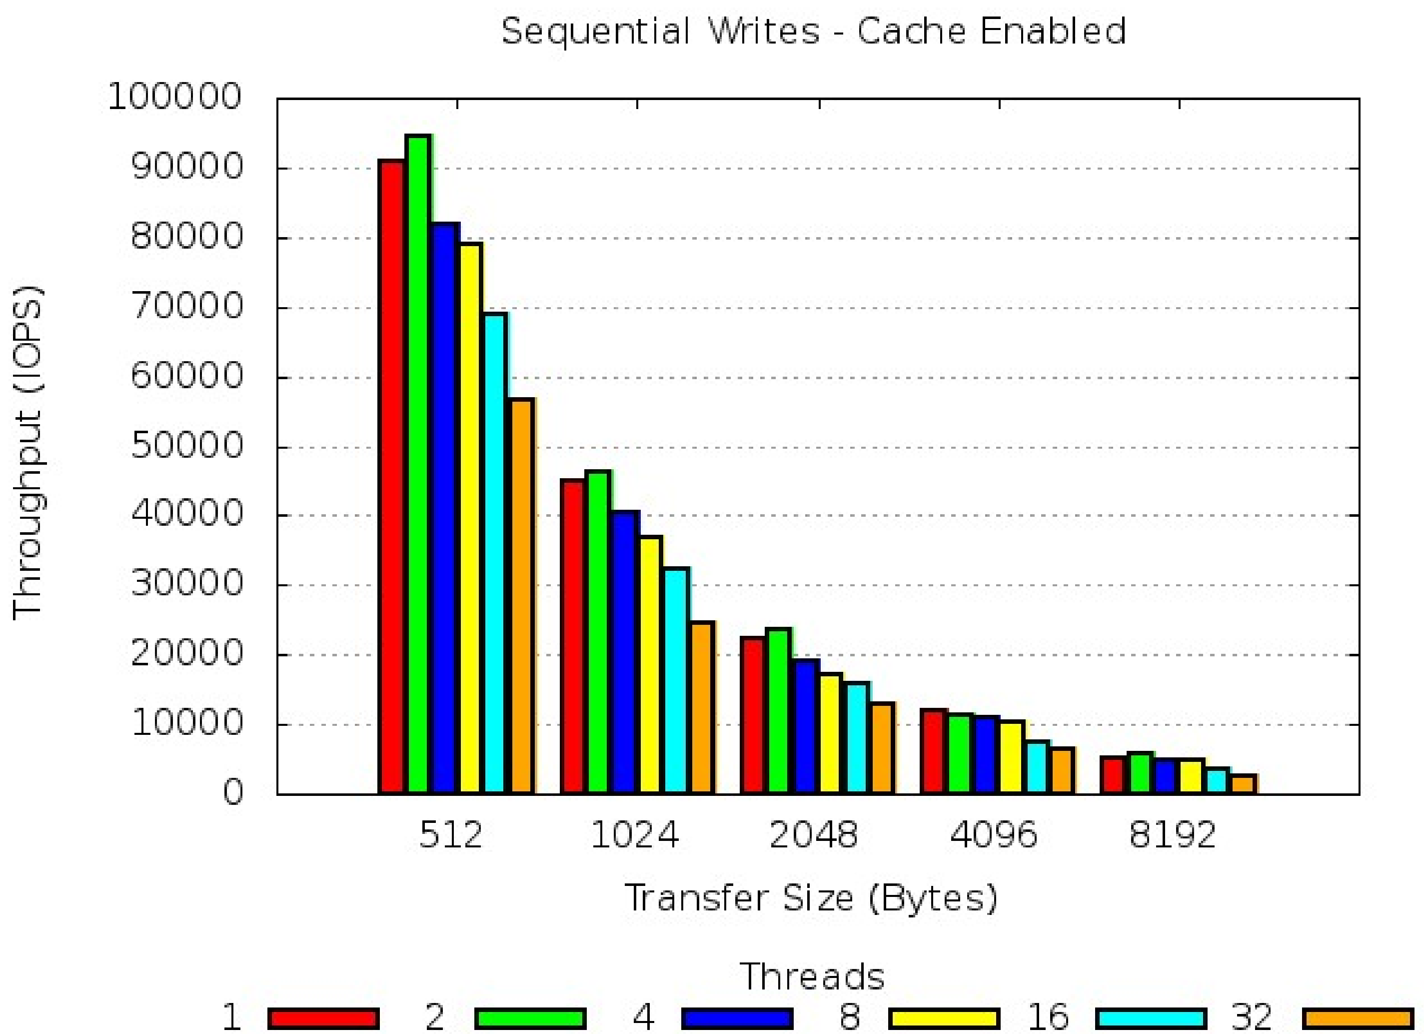

This diagram describes how the sustained, sequential ARCHIVE data throughput rate (IO's/sec) was affected as the file transfer size of the objects being written to the DR550 increased and the number of ARCHIVE application threads were also successively increased.

The File System Gateway cache was enabled, allowing writes to leverage the cache in the File System Gateway to achieve greater throughput when archiving data to the DR550.

The shape of the curve in the diagram shows the expected effect of increasing file transfer sizes causing a successive decrease in IO's/sec.

Summary

The IBM DR550 File System Gateway offers simple file-archiving capability for the DR550 without requiring application-specific enablement. It provides the customer with the additional flexibility to use the DR500 with the broad range of content management applications that support an NFS or CIFS protocol interface for their storage archives and retrievals.

This Performance Evaluation has demonstrated the highest levels of sustained ARCHIVE (writes) and RETRIEVE (reads) function throughput that the DR550 File System Gateway could sustain under different workload conditions using different numbers of concurrent application threads, both with and without FSG cache.

Acknowledgments

Many thanks to the SAN HW and Server Proven Test Department and their Manager, Carlos Lujan, who allowed us to use their equipment and time to generate the measurements which made this paper possible.

Thanks in particular to Dave DeHaan who provided us with valuable advice, technical consulting and DR550 hardware support, despite his own busy schedule.

Finally, we'd like to thank Mark Bayus, our DR550 DCT Program Manager, who gave encouragement and support for the Performance effort.