IBM
May 2021

IBM LTO 9 Tape Drive
RAO Performance Position Paper

Authors:
Tsuyoshi Miyamura (miyamuu@jp.ibm.com)
Osamu Matsumiya (MOSAMU@jp.ibm.com)
Tape systems development
Tokyo, Japan

Table of Contents

Introduction

The purpose of this position paper is to examine the Recommended Access Order (RAO) performance of the IBM LTO 9 Tape Drive (LTO 9 tape drive) in an open systems environment.

The operation and implementation of LTO RAO is described in the IBM LTO SCSI Reference (GA32-0928-04).

Results from similar runs for TS1160 RAO have been included for purposes of comparison.

LTO 9 Tape Drive Overview

The ninth generation IBM LTO 9 tape drive offers a great storage capacity and an excellent performance using technology designed for the mid-range open systems environment that include IBM Power Systems™; selected Oracle and Hewlett Packard servers and Intel servers running supported versions of Microsoft Windows or Linux.

There are 2 interfaces available for the LTO 9 tape drive:

- 8 Gbps Fibre Channel (FC-8)
- 12 Gbps SAS

The Fibre Channel interface and the SAS interface are available on full height models.

The IBM LTO 9 tape drive offers a native data rate of up to 400 MB/s an increase of 11% over the previous IBM LTO 8 tape drive.

The IBM LTO 9 tape drive supports a new generation of data cartridge (gen9 media) that offers a native capacity of 18000GB (45000GB with 2.5:1 compression) 1.5x the capacity of the previous gen8 tape cartridge generation.

The IBM LTO 9 tape drive will support the Linear Tape File System (LTFS) format in IBM Spectrum Archive™ that presents the tape storage as a file-based storage system. Additionally, the IBM LTO 9 tape drive is able to read and write previous generation 8 media to help customers protect their existing tape investments.

Performance Overview

The key features of the IBM LTO 9 tape drive are designed to improve performance and capabilities when compared to the IBM LTO 8, IBM LTO 7 and other vendors tape drives, some of the improvements are:

The LTO 9 tape drive features hardware encryption of data and two interface options. The data rate improvement from the previous generation is 11% and the tape capacity also increases 1.5x from the previous generation 8.

- Native data rate of up to 400 MB/s
- Native data physical capacity of 18000GB
- The data compression keeps the ratio to 2.5:1
- Support for 8Gb FC and 12Gb SAS connectivity
- SkipSync Function to provide small file backhitchless flush capability
- Cache buffer: 1024MB

This position paper examines the RAO performance benchmarks of the IBM LTO 9 tape drive.

Performance Evaluation

All of the performance benchmarks were run on the following systems:

- IBM System x3550 M5 server running RHEL 7.5 with QLogic ISP8324-based 16Gb Fibre Channel and N2225 12Gb SAS External HBA.

The performance benchmarks used for the tests are a toolbox of in-house C-based performance measurement tools designed to fully exercise the host interface and tape drive with the least amount of overhead. As such, the primary goal of the benchmarks was to provide a picture of the maximum capabilities of the LTO 9 tape drive. All data rates/capacity reflect a decimal basis where MB = 1,000,000 bytes and GB=1,000 MB. Actual tape drive data rate and cartridge capacity might vary depending on factors such as data compression, server and disk performance variables.

There are a number of factors that impact performance, especially data rate at high compression ratios and large block sizes. Server hardware performance, server slot and operating system/device driver performance are important factors. Another source of variability in the data rate performance tests could be due to the firmware used for the Fibre Channel and SAS connections.

All RAO tests were performed with tapes fully written with non-compressible data. The RAO implementation in LTO produces the best results for performance enhancement when there is little variability in block size or data compression ratio. When the variability in compression ratio or block sizes increase, the accuracy of the locate estimates may be reduced and any potential performance enhancements may be diminished.

Multiple RAO runs were performed using random UDS locations for each number of UDSes of the size indicated to show an average time. Some charts show the minimum and maximum times of these runs to illustrate potential variations in results.

LTO 9 L9 media RAO Performance

The Locate and Locate+Read performance of the IBM LTO 9 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating or locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

It took less time to locate and locate+read when the order of access was determined by RAO compared to when the order of access was determined by sorting by record number.

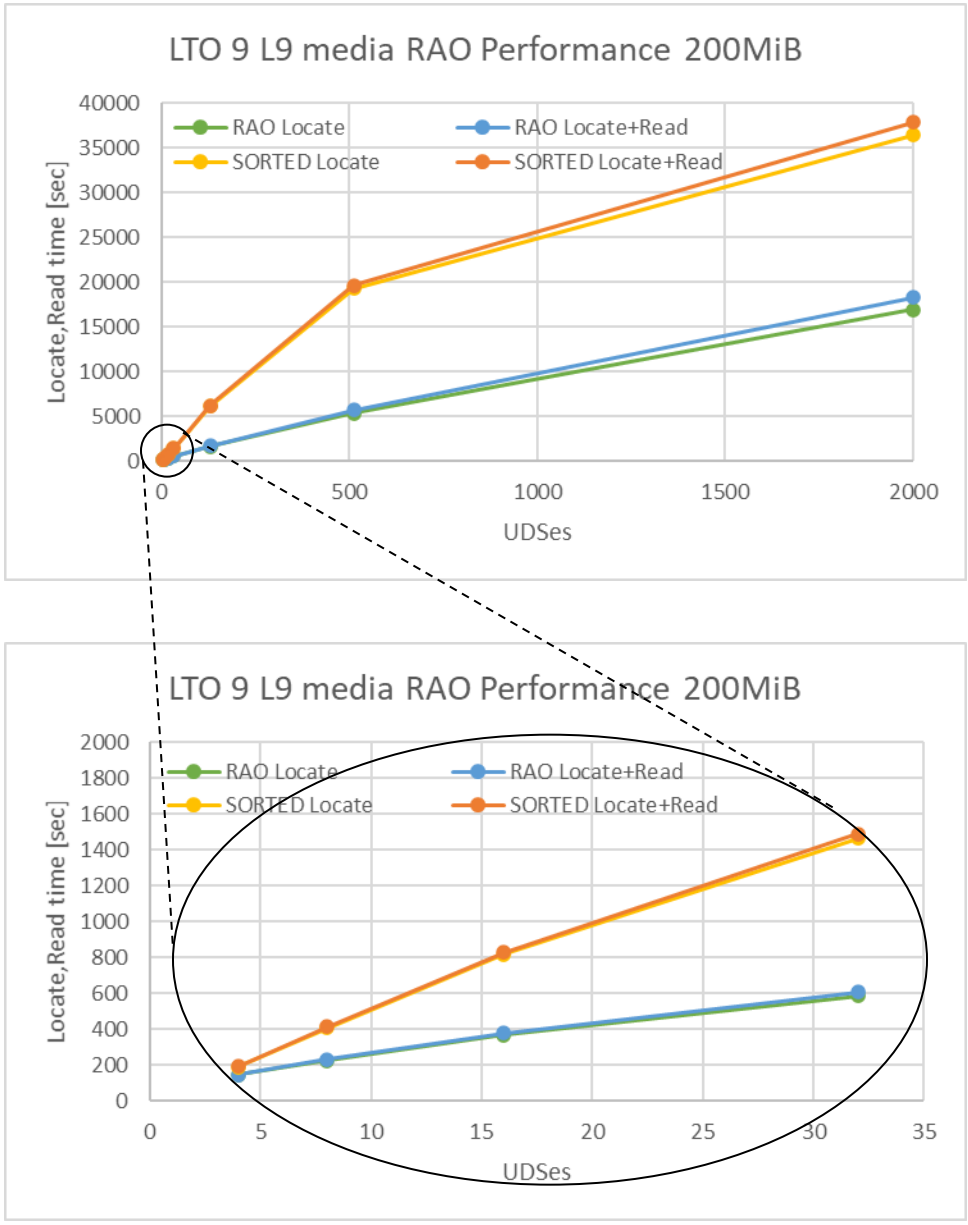

The Locate and Locate+Read performance of the IBM LTO 9 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating or locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

It took less time to locate and locate+read when the order of access was determined by RAO compared to when the order of access was determined by sorting by record number.

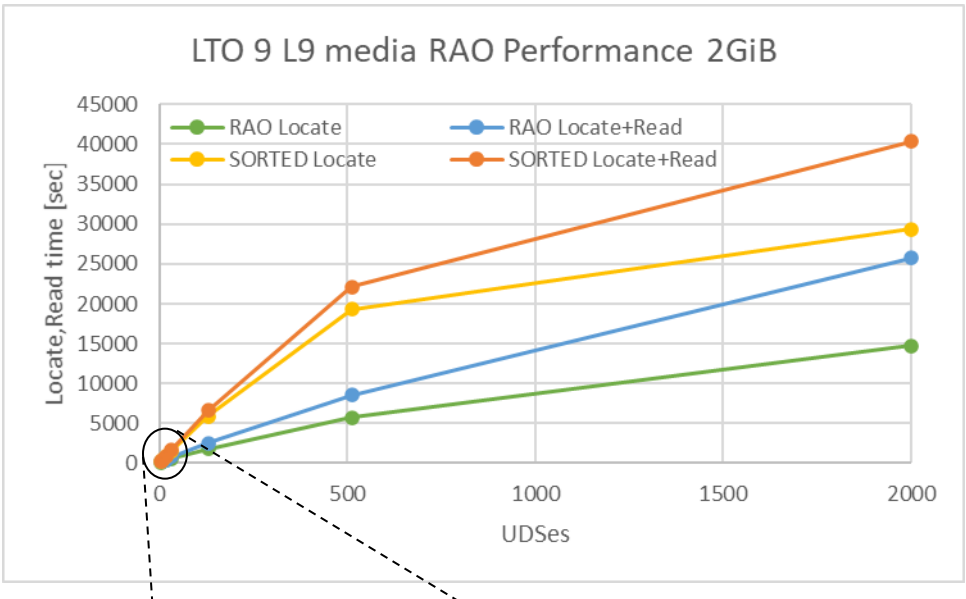

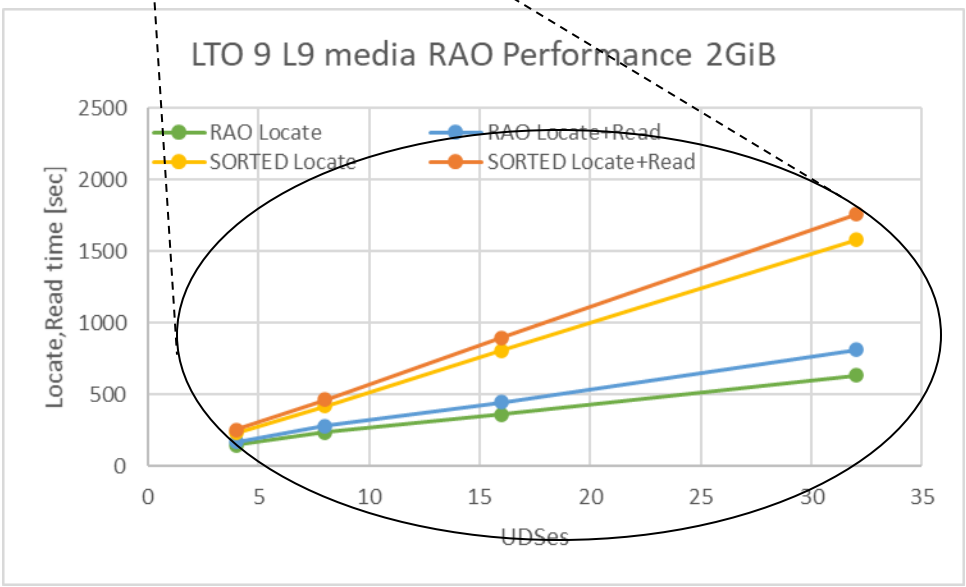

LTO 9 L9 media RAO Locate+Read Performance Range

The range of the Locate+Read performance of the IBM LTO 9 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

The larger the number of UDSes, the smaller the range of Locate+Read time. The max and min bars are almost overlapped on the average values (dot marks).

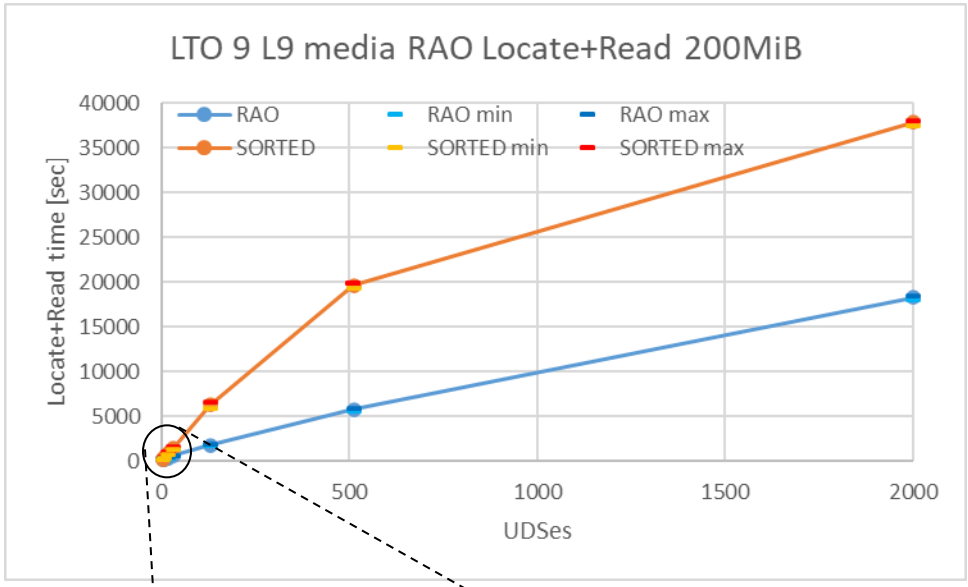

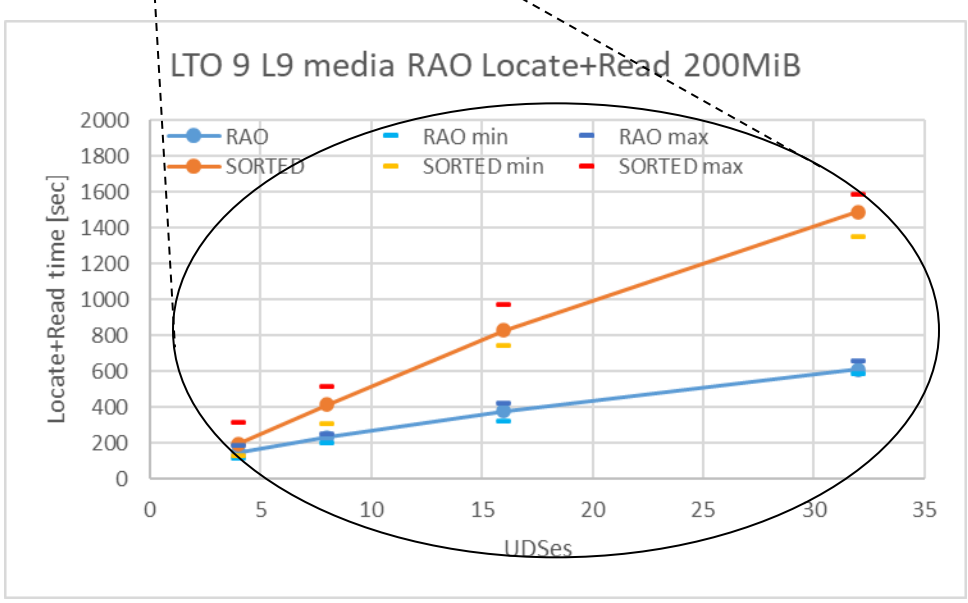

The range of the Locate+Read performance of the IBM LTO 9 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

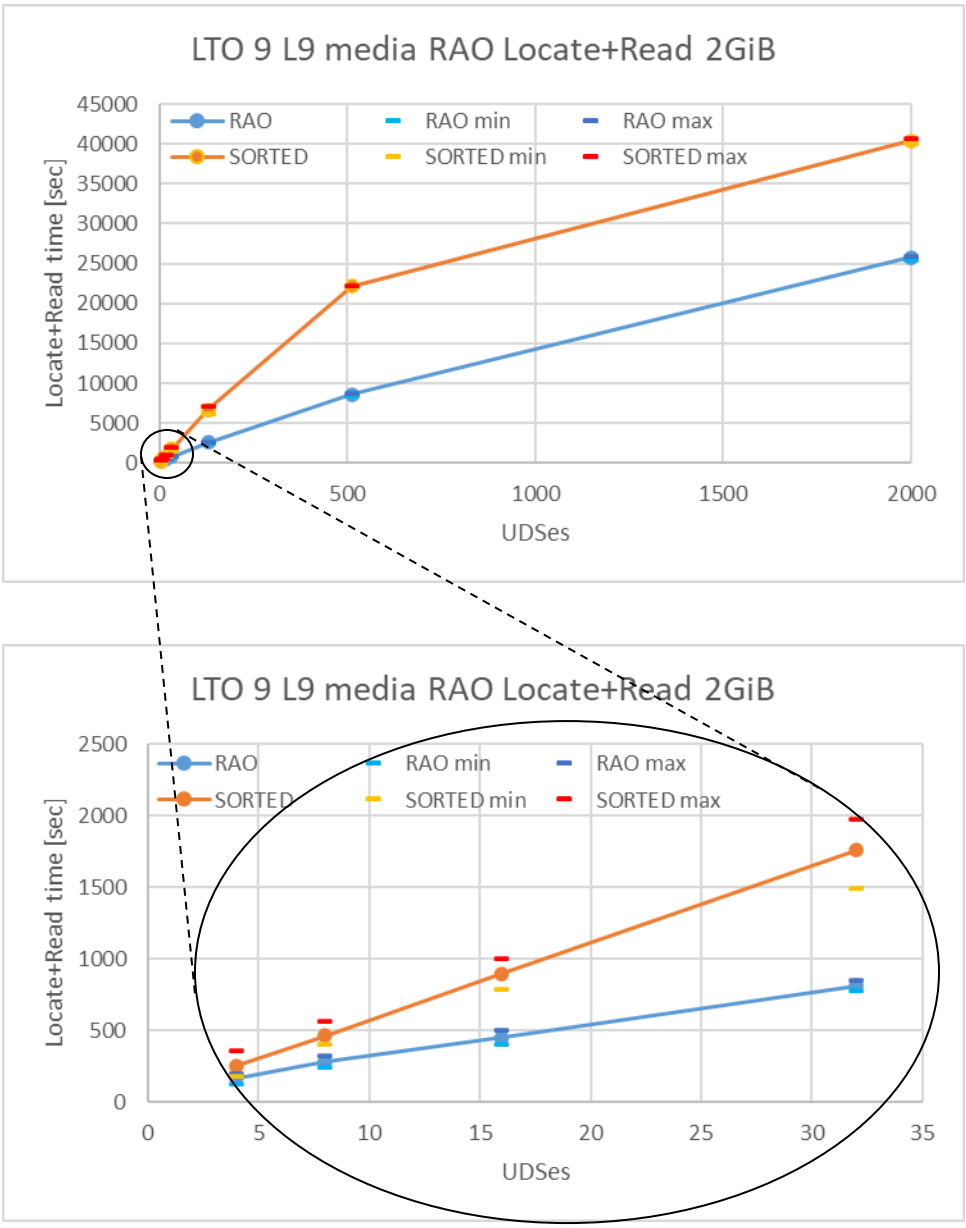

LTO 9 L9 media RAO Improvement versus Sorted

The following chart shows the Locate and Locate+Read time improvement for LTO 9 media on LTO 9 tape drive. The improvement is calculated as ((sorted - rao) / sorted) for both Locate and Locate+Read using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).
Outline of Locate improvement is colored in green with dot in yellow which expresses the RAO (green) Locate improvement versus Sorted (yellow).
Outline of Locate+Read improvement is colored in blue with dot in orange which expresses the RAO (blue) Locate+Read improvement versus Sorted (orange).

RAO's performance improvement over Sorted is maximized when the number of UDSes is 128.

In general, the smaller the read, more of the overall locate+read time is taken by the locate operations.

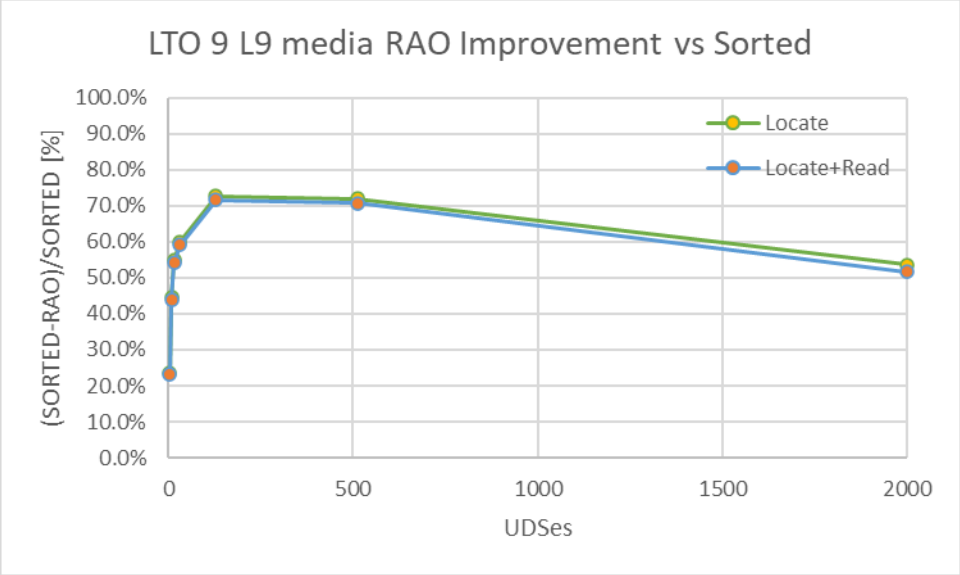

The following chart shows the Locate and Locate+Read time improvement for LTO 9 media on LTO 9 tape drive. The improvement is calculated by dividing Locate and Locate+Read time using Sorted order minus Locate and Locate+Read time using RAO by RAO and Sorted by Locate and Locate+Read time using Sorted order using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

In general, the larger the read, less of the overall locate+read time is taken by the locate operations.

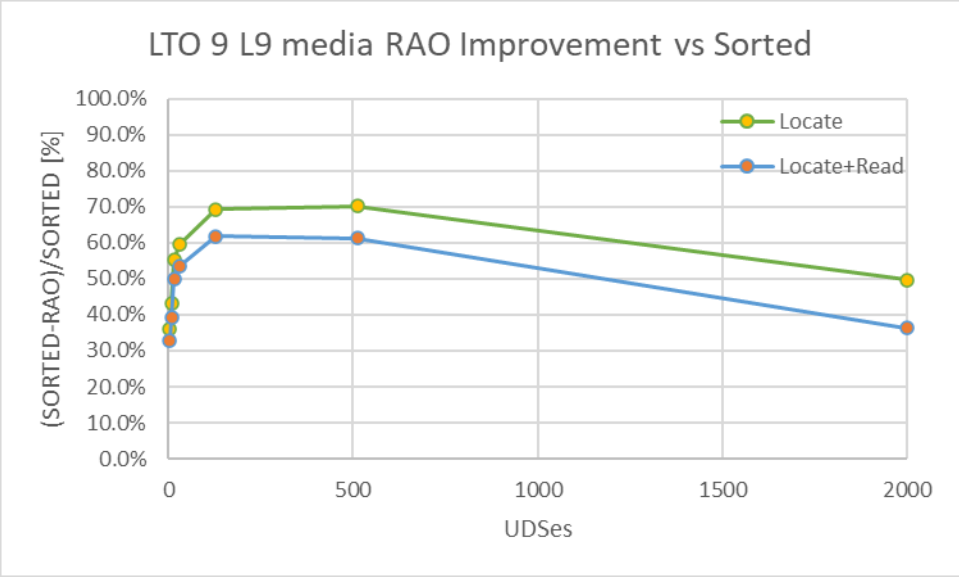

LTO 9 L8 media RAO Performance

The Locate and Locate+Read performance of the IBM LTO 8 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating or locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

It took less time to locate and locate+read when the order of access was determined by RAO compared to when the order of access was determined by sorting by record number.

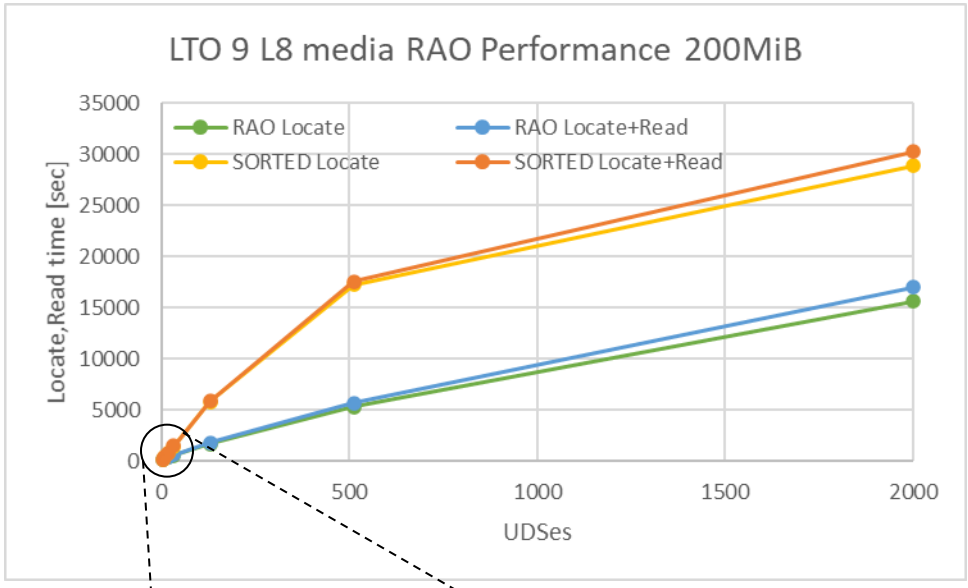

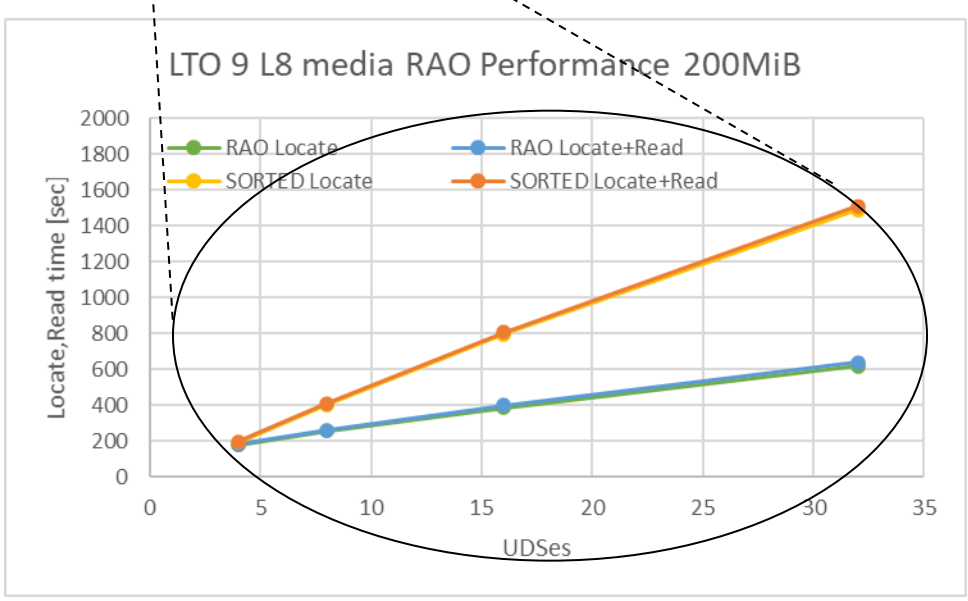

The Locate and Locate+Read performance of the IBM LTO 8 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating or locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

It took less time to locate and locate+read when the order of access was determined by RAO compared to when the order of access was determined by sorting by record number.

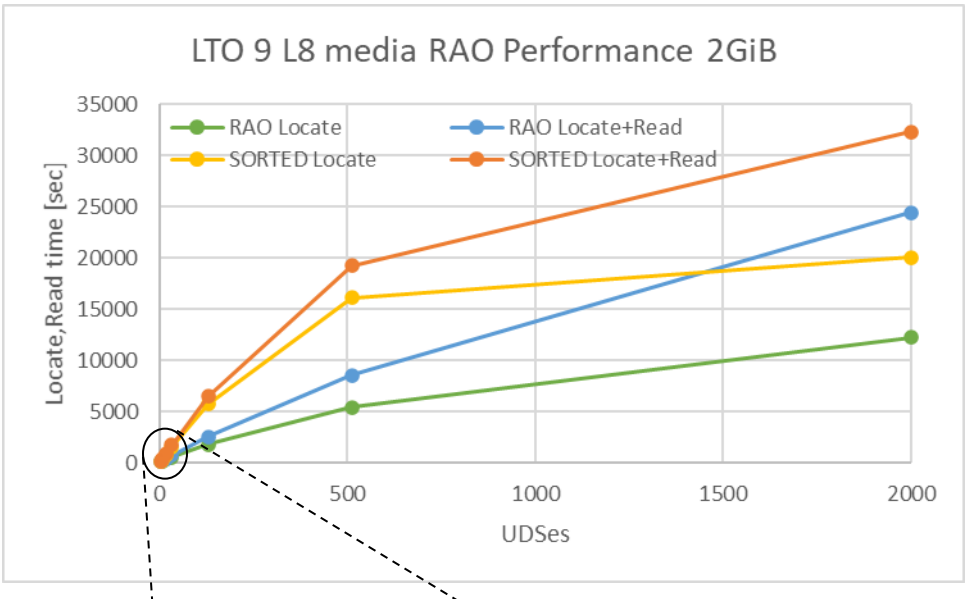

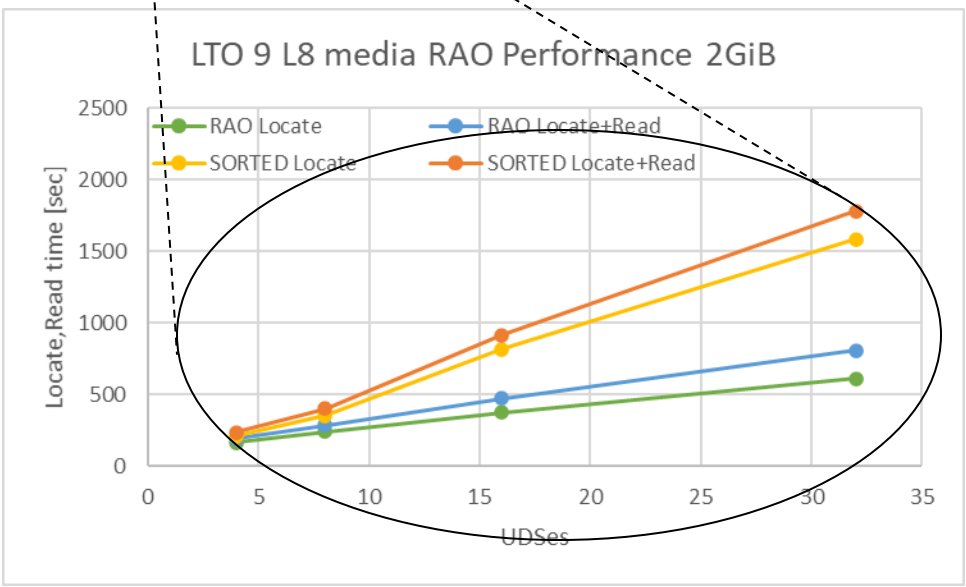

LTO 9 L8 media RAO Locate+Read Performance Range

The range of the Locate+Read performance of the IBM LTO 8 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

The larger the number of UDSes, the smaller the range of Locate+Read time.

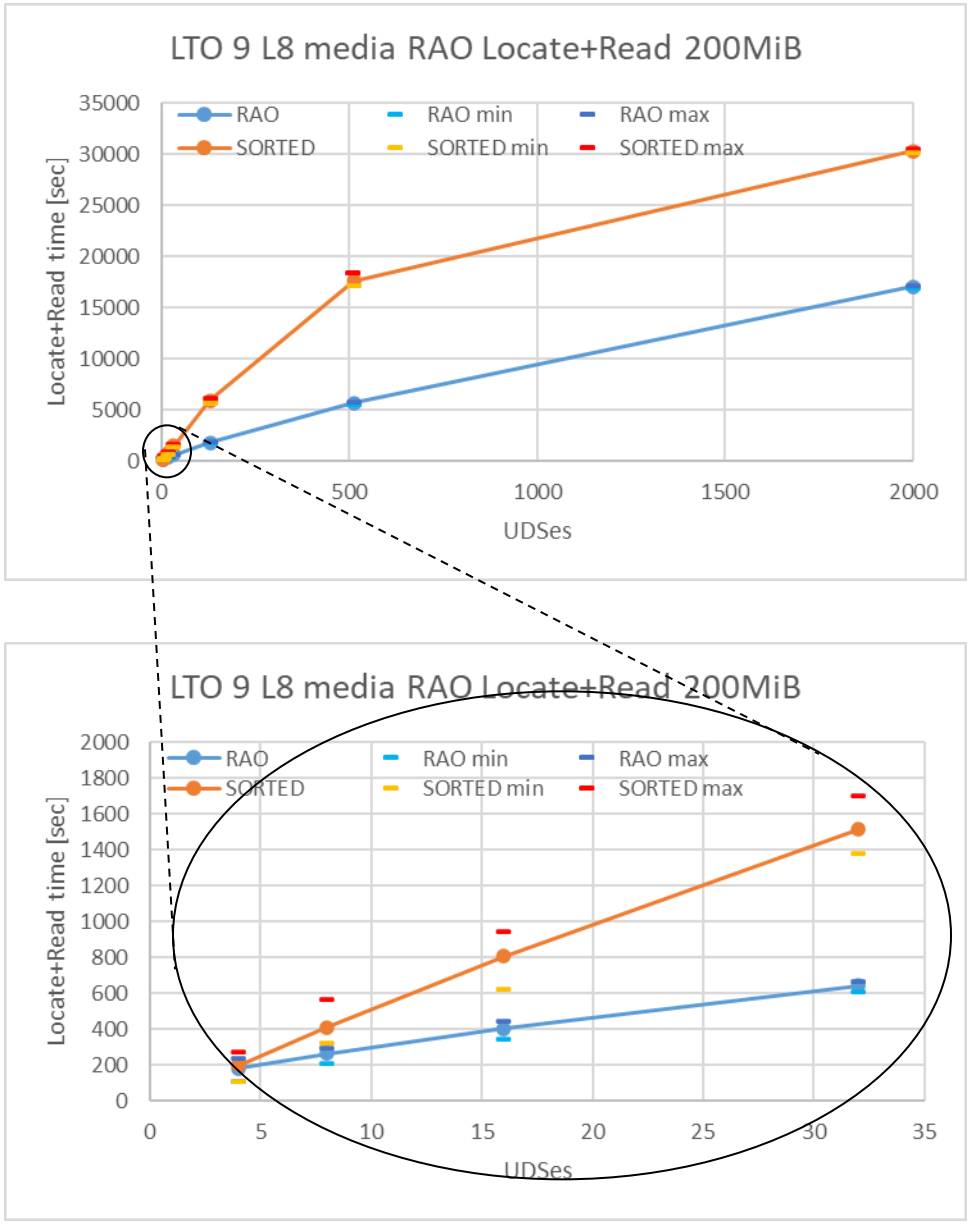

The range of the Locate+Read performance of the IBM LTO 8 tape media on the IBM LTO 9 tape drive is described by the following set of charts that show how the drive behaves when locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

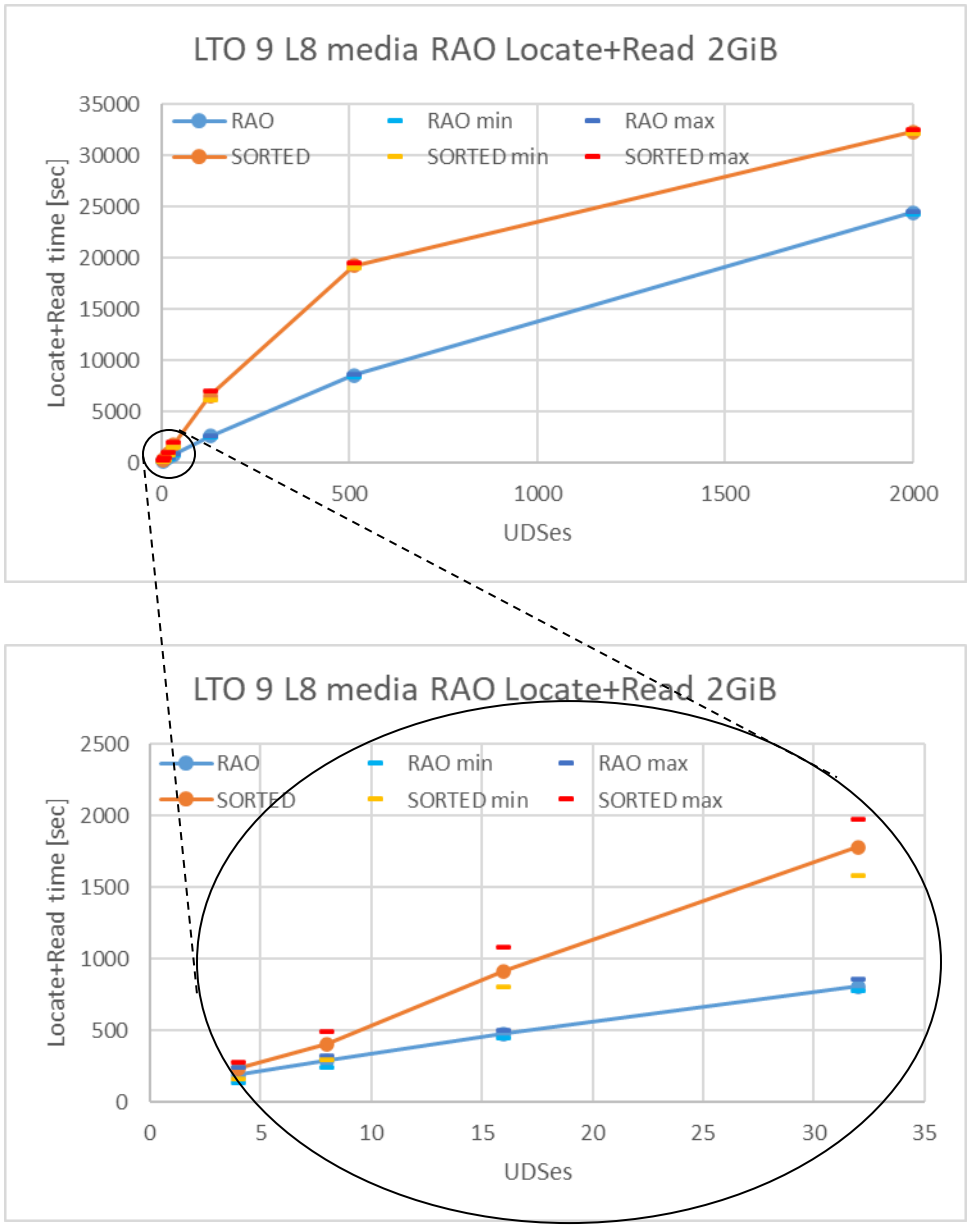

LTO 9 L8 media RAO Improvement versus Sorted

The following chart shows the Locate and Locate+Read time improvement for LTO 8 media on LTO 9 tape drive. The improvement is calculated as ((sorted - rao) / sorted) for both Locate and Locate+Read using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

RAO's performance improvement over Sorted is maximized when the number of UDSes is 128.

In general, the smaller the read, more of the overall locate+read time is taken by the locate operations.

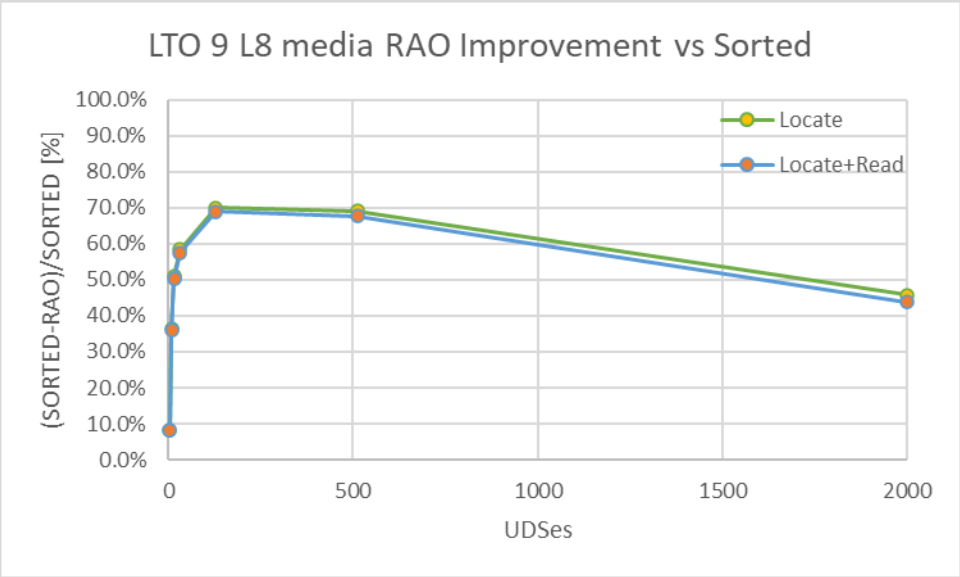

The following chart shows the Locate and Locate+Read time improvement for LTO 8 media on LTO 9 tape drive. The improvement is calculated by dividing Locate and Locate+Read time using Sorted order minus Locate and Locate+Read time using RAO by RAO and Sorted by Locate and Locate+Read time using Sorted order using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

In general, the larger the read, less of the overall locate+read time is taken by the locate operations.

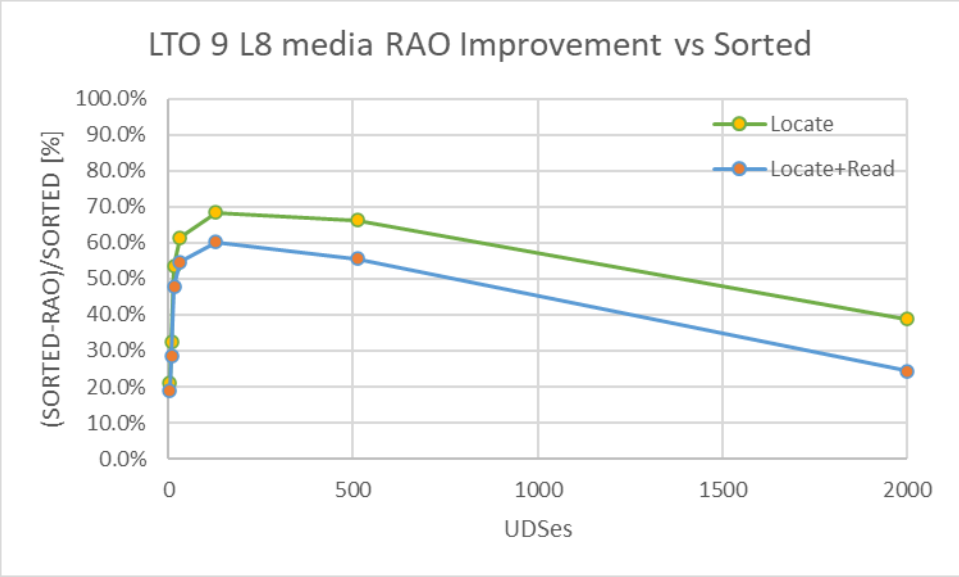

TS1160 JE media RAO Performance

To compare the performance with TS1160, the Locate and Locate+Read performance of the JE tape media on the IBM TS1160 tape drive is described by the following set of charts that show how the drive behaves when locating or locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

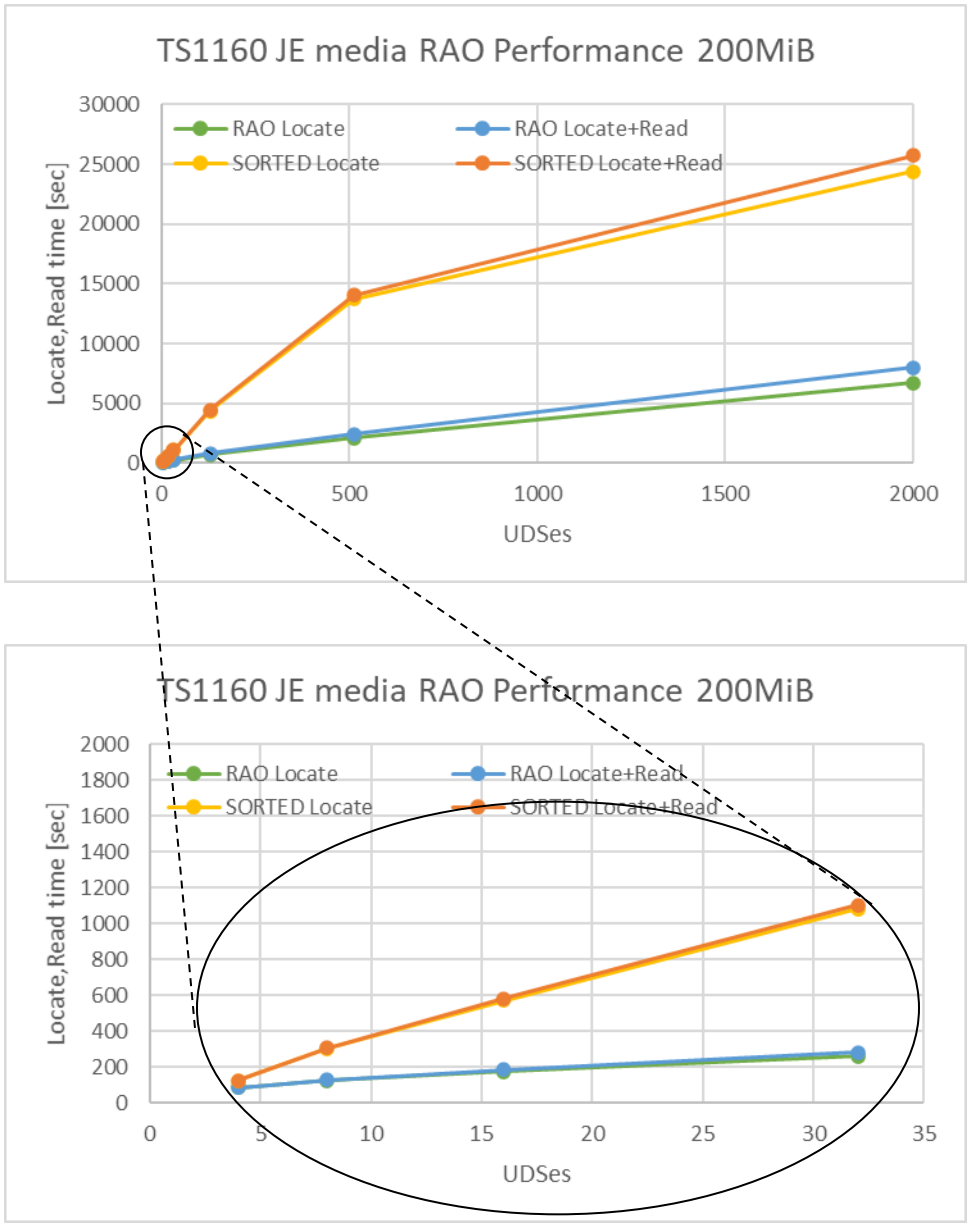

The Locate and Locate+Read performance of the JE tape media on the IBM TS1160 tape drive is described by the following set of charts that show how the drive behaves when locating or locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

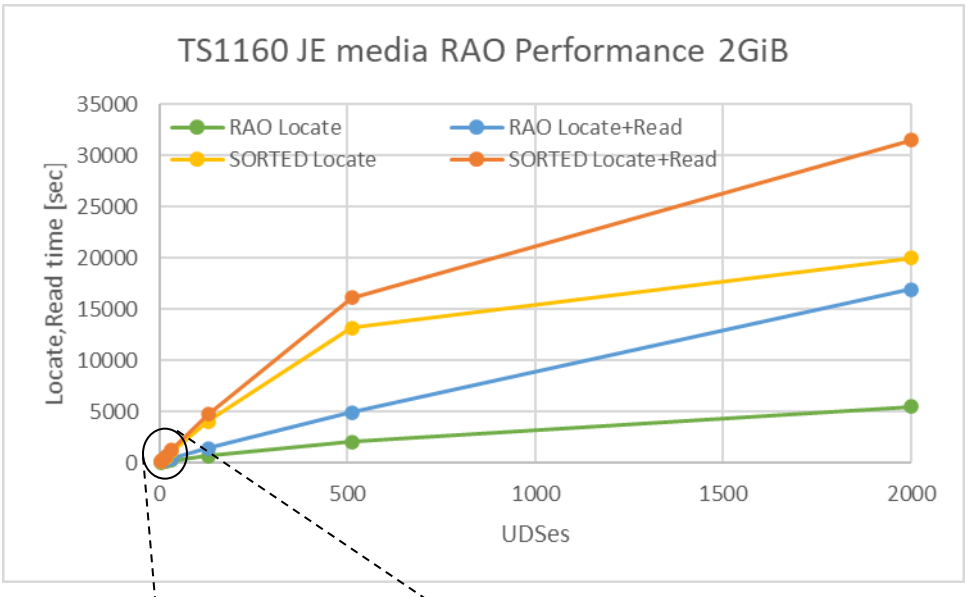

It took less time to locate and locate+read when the order of access was determined by RAO compared to when the order of access was determined by sorting by record number.

RAO performance using JE tape on TS1160 is better than LTO 9 drive because TS1160 utilizes HRTD mechanism to improve locate performance.

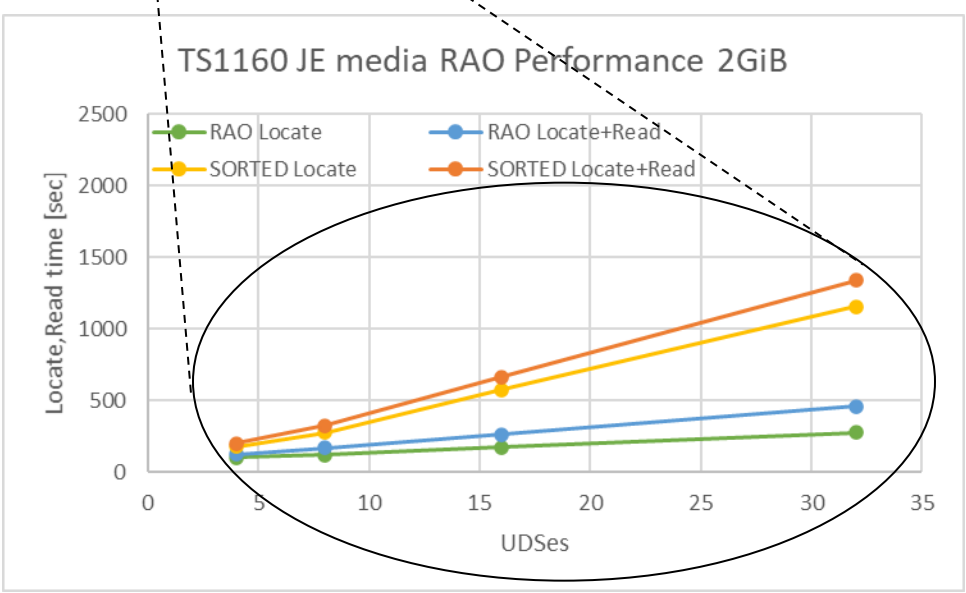

TS1160 JE media RAO Locate+Read Performance Range

The range of the Locate+Read performance of the JE tape media on the IBM TS1160 tape drive is described by the following set of charts that show how the drive behaves when locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

The larger the number of UDSes, the smaller the range of Locate+Read time.

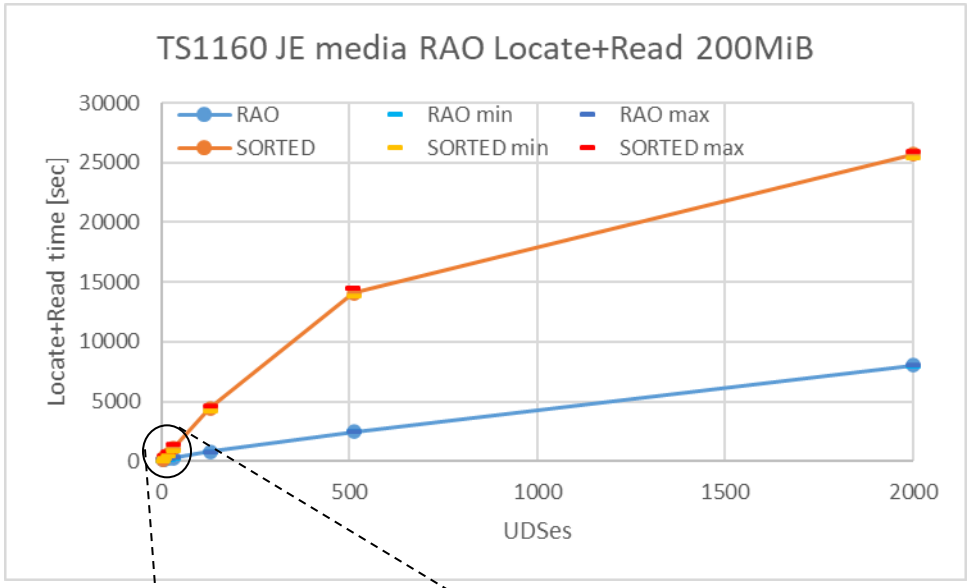

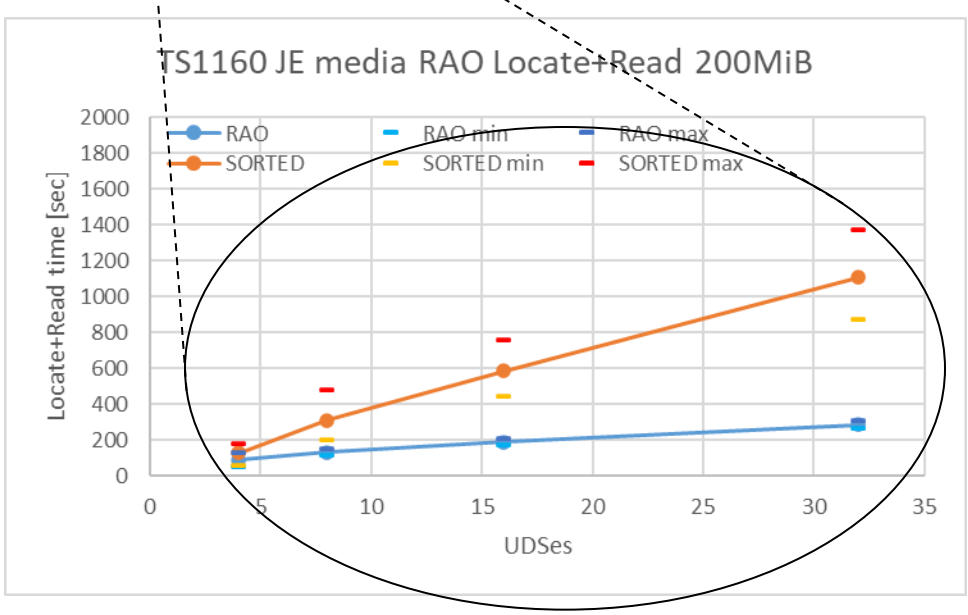

The range of the Locate+Read performance of the JE tape media on the IBM TS1160 tape drive is described by the following set of charts that show how the drive behaves when locating+reading data in the access order determined by RAO and Sorted by record number using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

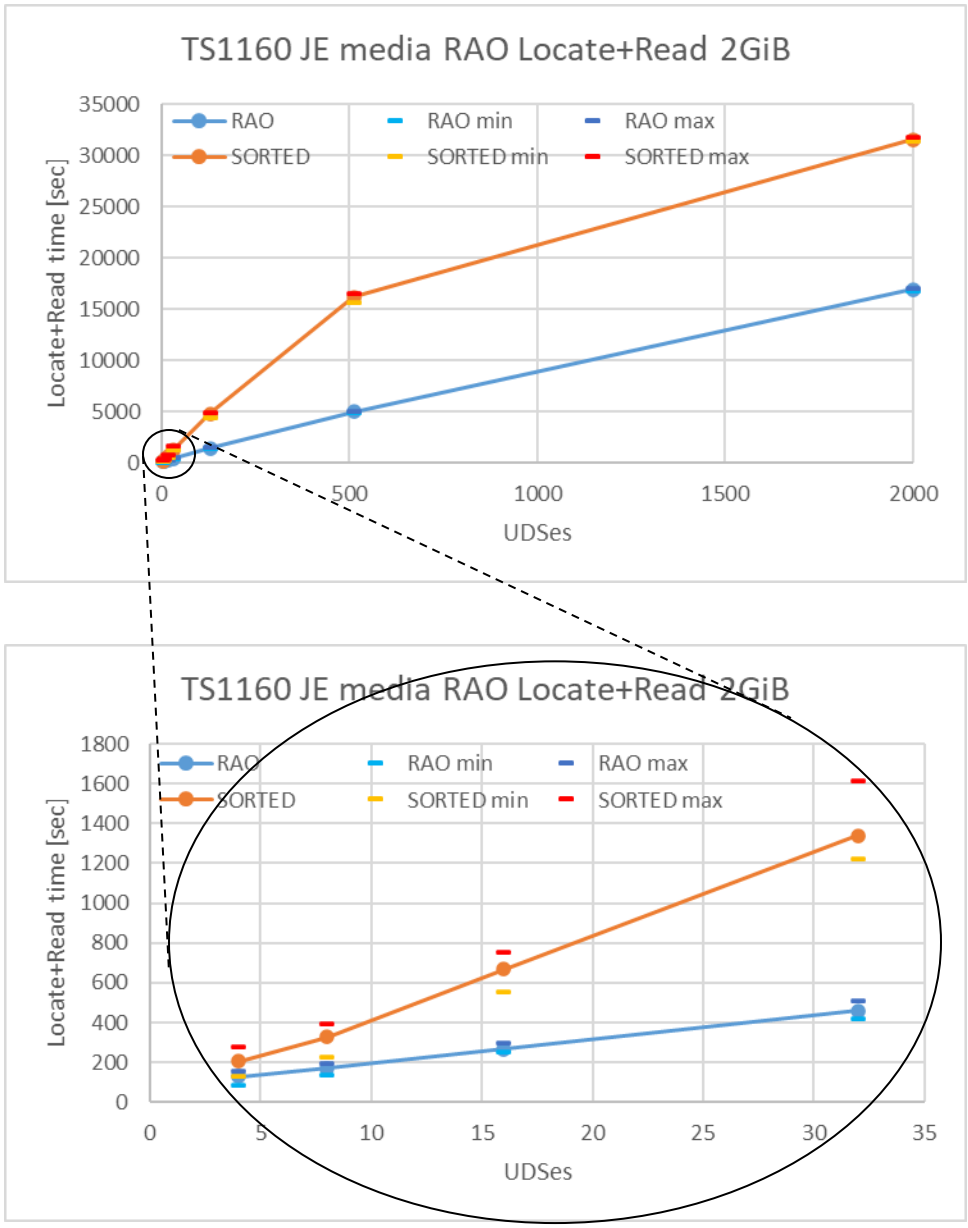

TS1160 JE media RAO Improvement versus Sorted

The following chart shows the Locate and Locate+Read time improvement for JE media on TS1160 tape drive. The improvement is calculated as ((sorted - rao) / sorted) for both Locate and Locate+Read using differing numbers of 200MiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

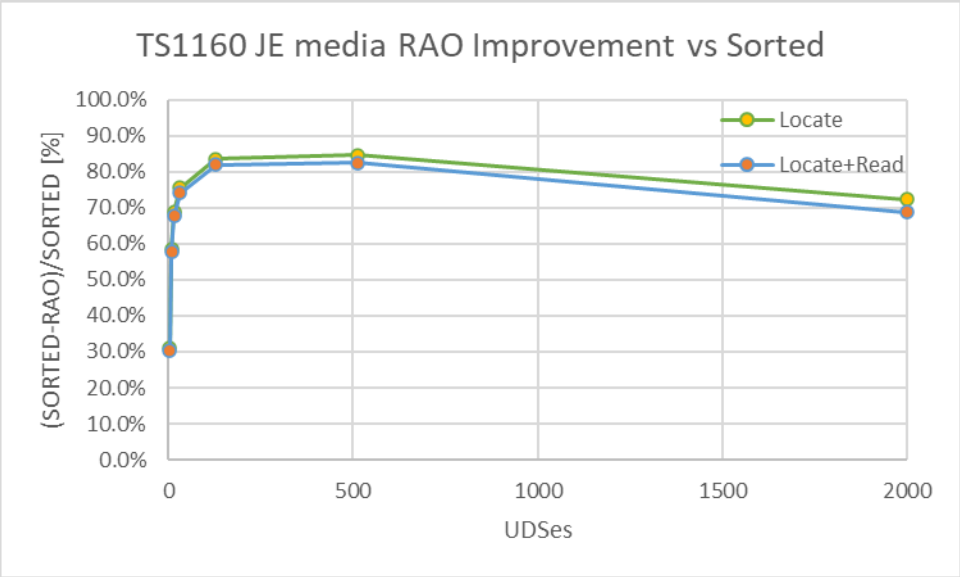

RAO's performance improvement over Sorted is maximized when the number of UDSes is 128.

In general, the smaller the read, more of the overall locate+read time is taken by the locate operations.

The following chart shows the Locate and Locate+Read time improvement for JE media on TS1160 tape drive. The improvement is calculated by dividing Locate and Locate+Read time using Sorted order minus Locate and Locate+Read time using RAO by RAO and Sorted by Locate and Locate+Read time using Sorted order using differing numbers of 2GiB UDSes (4, 8, 16, 32, 128, 512 and 2000).

In general, the larger the read, less of the overall locate+read time is taken by the locate operations.

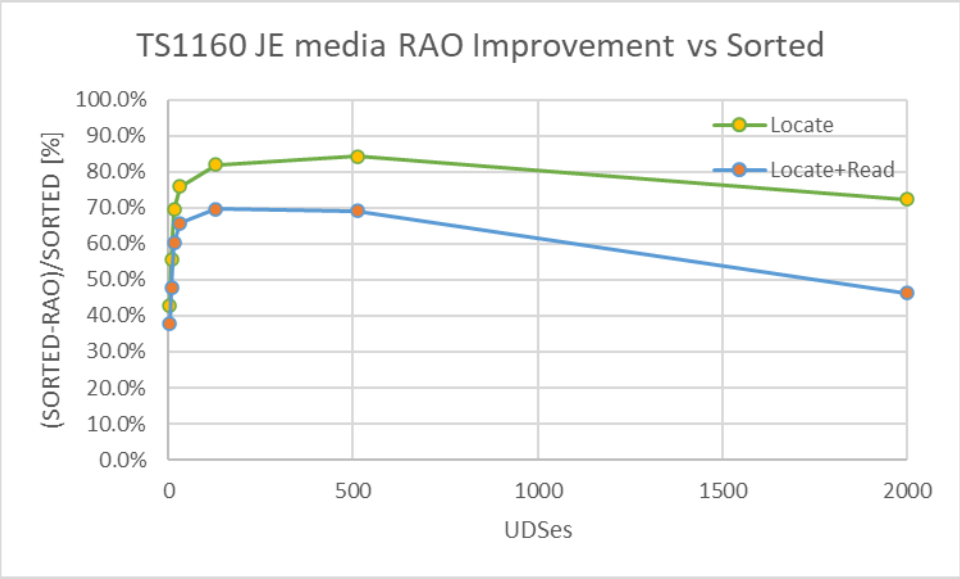

Conclusions

Since the introduction of the first LTO tape drive, every following generation has incorporated new features and performance improvements to respond to storage needs. Now the IBM LTO 9 tape drives in conjunction with the new LTO Gen 9 media represent an efficient solution for today's growing storage demands.

Locate time and Locate+Read time is improved by using access order calculated by RAO.

In general, the smaller the read, more of the overall locate+read time is taken by the locate operations.

The IBM LTO 9 tape drive is a smart storage solution for businesses requiring backup and archival storage of their data.

This point is visually indicated in the Improvment charts. Locate time and Locate+Read time is improved by using access order calculated by RAO. When reads are relatively small, most of the overall time is spent locating. As a result, RAO provides the highest percentage improvement in such cases as shown in the RAO Improvement vs Sorted charts. However, absolute time improvement when using RAO is primarily based on the number of UDSes (read sizes are typically not a significant factor).

International Business Machines Corporation 2021
IBM Systems Group
5600 Cottle Road,
San Jose, CA 95193

Printed in the United States of America
May 2021

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 and/or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice.

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 IBM's intellectually 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 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.