IBM TotalStorage Virtual Tape Server

August 30, 2002

FICON and Advanced Policy Management Performance Overview - Version 1.0

J.D. Metzger & Jesse Thrall
Tucson, AZ
jdmetzg@us.ibm.com
thrall@us.ibm.com

Trademarks

The following are trademarks or registered trademarks of the International Business Machines Corporation:

IBM	OS/390	ESCON
S/390	FICON	zSeries
TotalStorage		

Other company, product, and service names may be trademarks or registered trademarks of their respective companies.

Disclaimers - Part 1

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein 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. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

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 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.

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.

Disclaimers - Part 2

The responsibility for use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's or user's ability to evaluate and integrate them into their operating environment. Customers or users attempting to adapt these techniques to their own environments do so at their own risk. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

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.

Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Contents

Introduction

Introduction

On August 30, 2002, IBM delivered significant enhancements to the capabilities of the Total Storage Virtual Tape Server:

- FICON attachment for models B10 and B20
- Advanced Policy Management, including
  - Physical Volume Pooling
  - Selective Dual Copy
  - Tape Volume Cache Management
  - Peer-to-Peer Selective Copy Mode

This document provides an overview of the performance impact of these enhancements as revealed in laboratory measurements.

VTS FICON Attachment Overview

VTS FICON Support

- VTS Models B10 & B20 only
  - Long and Short Wave attachment features
  - B20 FICON Performance Accelerator Feature (FPAF)
  - Channel Adapter Intermix ESCON / FICON
  - FICON Director Support
  - IBM 2042 Inrange FC/9000 Fibre Channel Director
  - IBM 2032 McData ED-6064 Enterprise Fibre Channel Director
  - Feature upgrades and conversions

Note: Peer-to-Peer VTS FICON support not announced

New VTS FICON Features

Feature	Description	Prerequisite	Model
0201	9 Micron LC/LC 31 meter cable	3415	B10, B20
0202	9 Micron LC/SC 31 meter cable	3415	B10, B20
0203	50 Micron LC/LC 31 meter cable	3416	B10, B20
0204	50 Micron LC/SC 31 meter cable	3416	B10, B20
0205	62.5 Micron LC/LC 31 meter cable	3416	B10, B20
0206	62.5 Micron LC/SC 31 meter cable	3416	B10, B20
3000	FICON Enablement	None	B10, B20
3415	FICON attach - Long Wavelength	3000	B10, B20
3416	FICON attach - Short Wavelength	3000	B10, B20
5250	FICON Performance Accelerator	3415/3416	B20
9700	No FICON cable from plant	None	B10,B20

VTS Connectivity Options

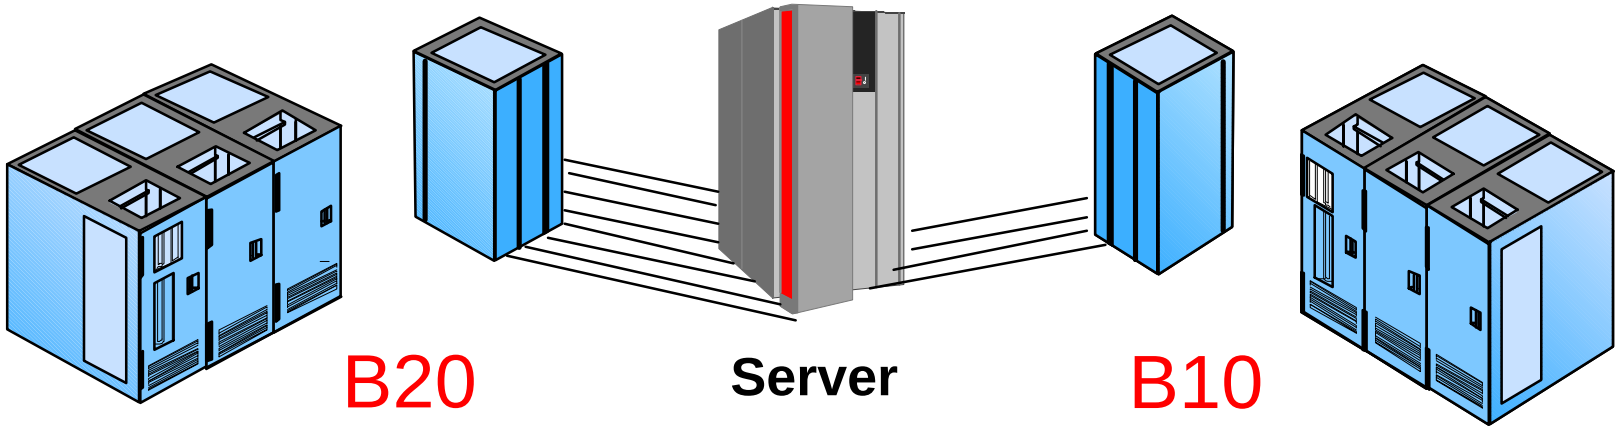

# ESCON	# FICON	# SCSI
16	0	0
8	4	0
8	0	8
8	0	0
4	4	0
0	4	0
0	8	0

# ESCON	# FICON	# SCSI
4	0	4
4	0	0
2	2	0
2	0	0
0	4	0
0	2	0
0	0	8

New options

VTS FICON Performance Overview

VTS FICON Performance - General Statements

- FICON attachment can significantly boost VTS performance. This boost will be seen both in single and multi-job performance.
- For single jobs, running a job over a FICON channel will result in faster job time than for over an ESCON channel.
- For multiple jobs, running over FICON channels will provide increased peak write and read hit data rate capabilities, and in some cases increased sustained data rate capabilities.
- Because of the greater bandwidth provided by FICON, one FICON channel can typically replace four existing ESCON channels while maintaining the same level of channel performance.

Definitions

Peak write throughput:

The average host write data rate when unconstrained by any other activity in the VTS other than possible concurrent premigration activity. This is the maximum throughput that can be expected from the VTS, given the configuration and data characteristics assumed.

Sustained write throughput:

The average host write data rate when constrained by the VTS to prevent the overfilling of cache. Sustained write is the maximum throughput that the VTS can achieve when the VTS needs to make room in the cache for new data as it is being written. This rate is determined basically by the rate at which the VTS can copy data from the cache to the physical drives.

VTS FICON Performance - Measurements

The following charts are divided into four categories:

- Single-job / single-channel sensitivity measurements
  (Charts 1 - 4)
- Multi-job / multi-channel model/feature comparative measurements
  (Charts 5 - 15)
- Multi-job / multi-channel sensitivity measurements
  (Charts 16 - 21)
- Single-job / single-channel distance measurements
  (Charts 22 - 25)

VTS FICON Performance Charts

- Single-job / single-channel sensitivity measurements
  - Chart 1: FICON Single-job Write Performance
  - Chart 2: FICON Single-job Read Performance
  - Chart 3: FICON Single-job BUFNO Performance
  - Chart 4: ESCON / FICON Single-job Blocksize Performance
- Multi-job / multi-channel model/feature comparative measurements
  - Chart 5: B10 ESCON / FICON Multi-job Write and Read Hit Performance
  - Chart 6: B20 ESCON / Base FICON Multi-job Peak Write Performance
  - Chart 7: B20 ESCON / PAF FICON Multi-job Peak Write Performance
  - Chart 8: B20 Base / PAF FICON Multi-job Peak Write Performance
  - Chart 9: B20 ESCON / FICON Multi-job Read Hit Performance
  - Chart 10: B20 ESCON / Base FICON Multi-job Sustained Write Performance
  - Chart 11: B20 ESCON / PAF FICON Multi-job Sustained Write Performance
  - Chart 12: B20 ESCON / Base FICON Multi-job Peak Performance Vs. # Physical Drives
  - Chart 13: B20 ESCON / Base FICON Multi-job Sustained Performance Vs. # Physical Drives
  - Chart 14: B20 ESCON / Base FICON Multi-job Peak Write Performance - BUFNO=5
  - Chart 15: B20 Base / PAF FICON Multi-job Peak Write Duration

VTS FICON Performance Charts (con't)

- Multi-job / multi-channel sensitivity measurements
  - Chart 16: B10 FICON Multi-job Write Performance Vs. Data Compression
  - Chart 17: B20 Base FICON Multi-job Write Performance Vs. Data Compression
  - Chart 18: B20 Base FICON Multi-job Write Performance Vs. # Virtual Drives
  - Chart 19: B20 Base FICON Multi-job Write Performance Vs. Blocksize
  - Chart 20: B20 Base FICON Multi-job Write Performance Vs. Volume Size
  - Chart 21: B20 Base FICON Multi-job Read Performance Vs. Data Compression
- Single-job / single-channel distance measurements
  - Chart 22: FICON Single-job Write Performance Vs. Distance Vs. Blocksize
  - Chart 23: FICON Single-job Read Performance Vs. Distance Vs. Blocksize
  - Chart 24: ESCON / FICON Single-job Write Performance Vs. Distance
  - Chart 25: ESCON / FICON Single -job Read Performance Vs. Distance

VTS FICON Performance Measurement Metrics

Unless otherwise specified on individual charts, the following "standard workload" job characteristics apply to all measurements:

- 800 MB logical volumes
- QSAM BUFNO = 20
- 32KB data blocksize
- Maximum VTS cache size
- Maximum number of host channels
- Maximum number of physical tape drives
- 2.66:1 data compression factor
- 64 virtual drives active for B10
- 128 virtual drives active for B20
- VTS at local distance to host processors

All Tucson laboratory FICON measurements were made on a G6 host processor.

Chart 1: FICON
Single-job Write Performance

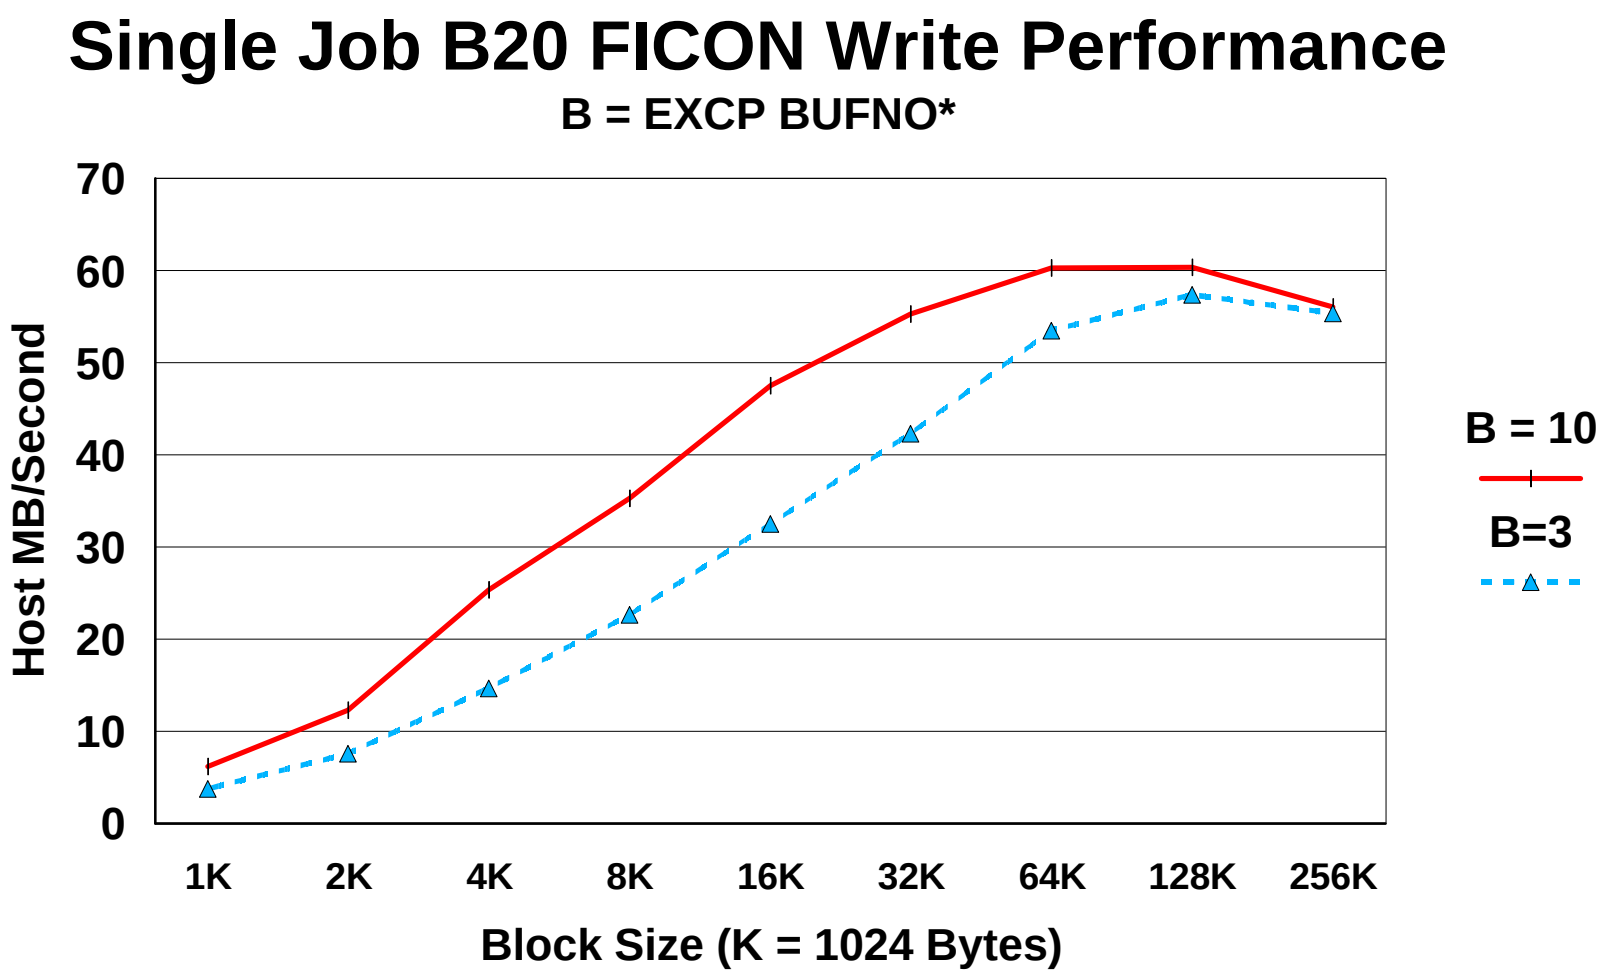

*** QSAM BUFNO = 2X EXCP BUFNO**

Chart 2: FICON
Single-job Read Performance

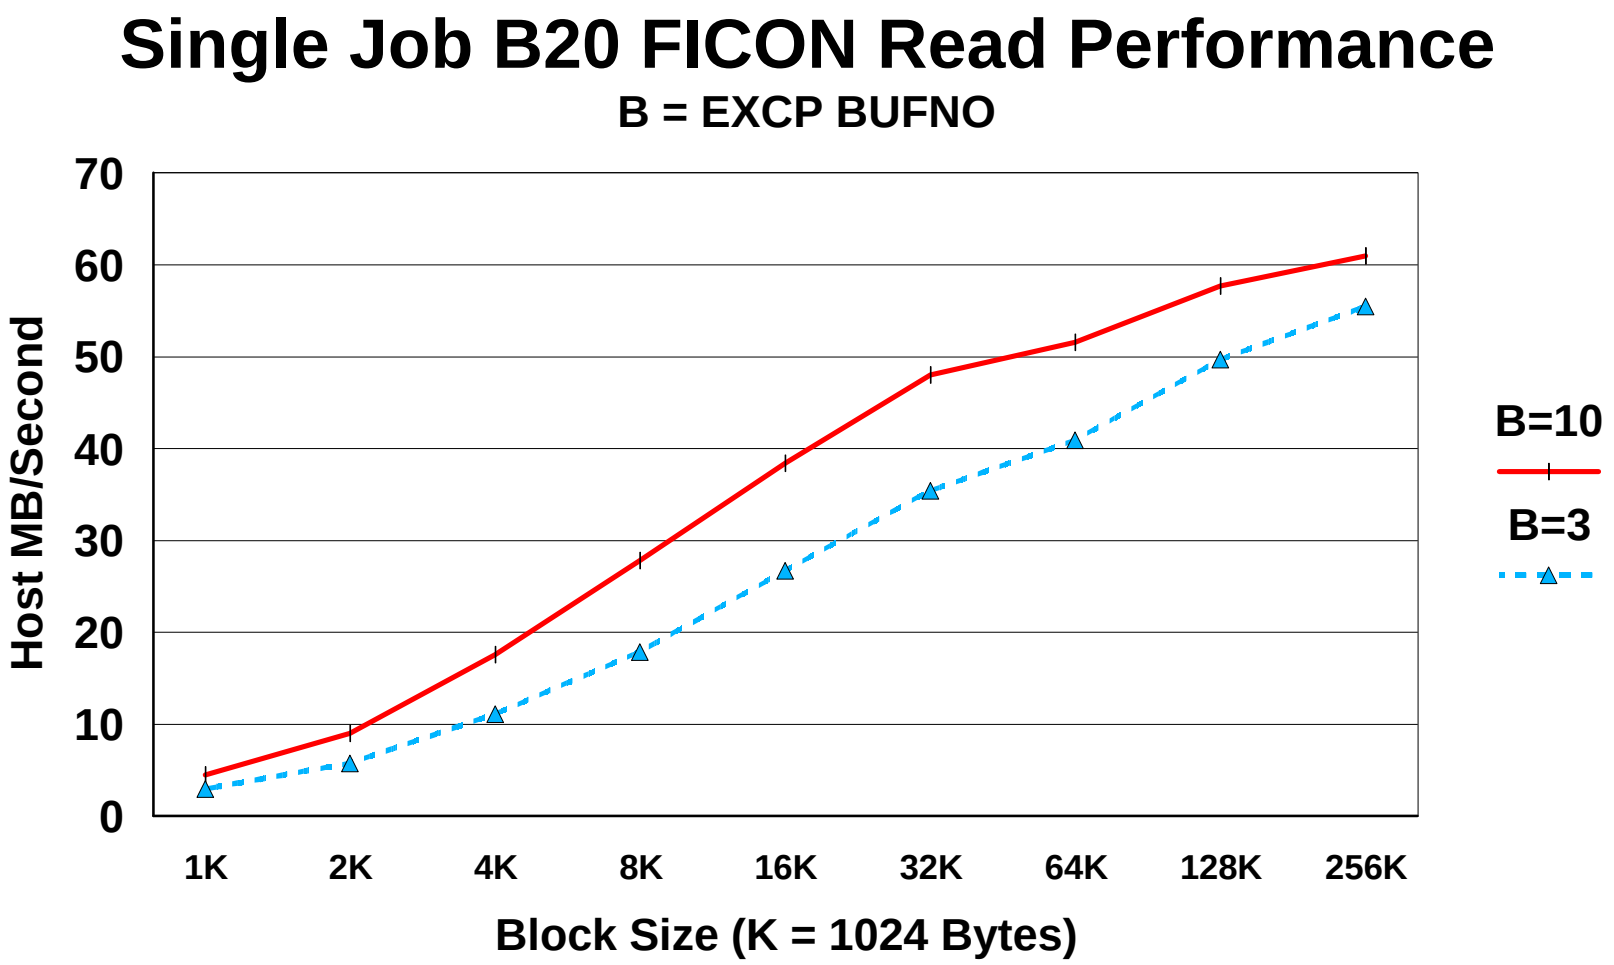

*** QSAM BUFNO = 2X EXCP BUFNO**

Chart 3: FICON
Single-job BUFNO Performance

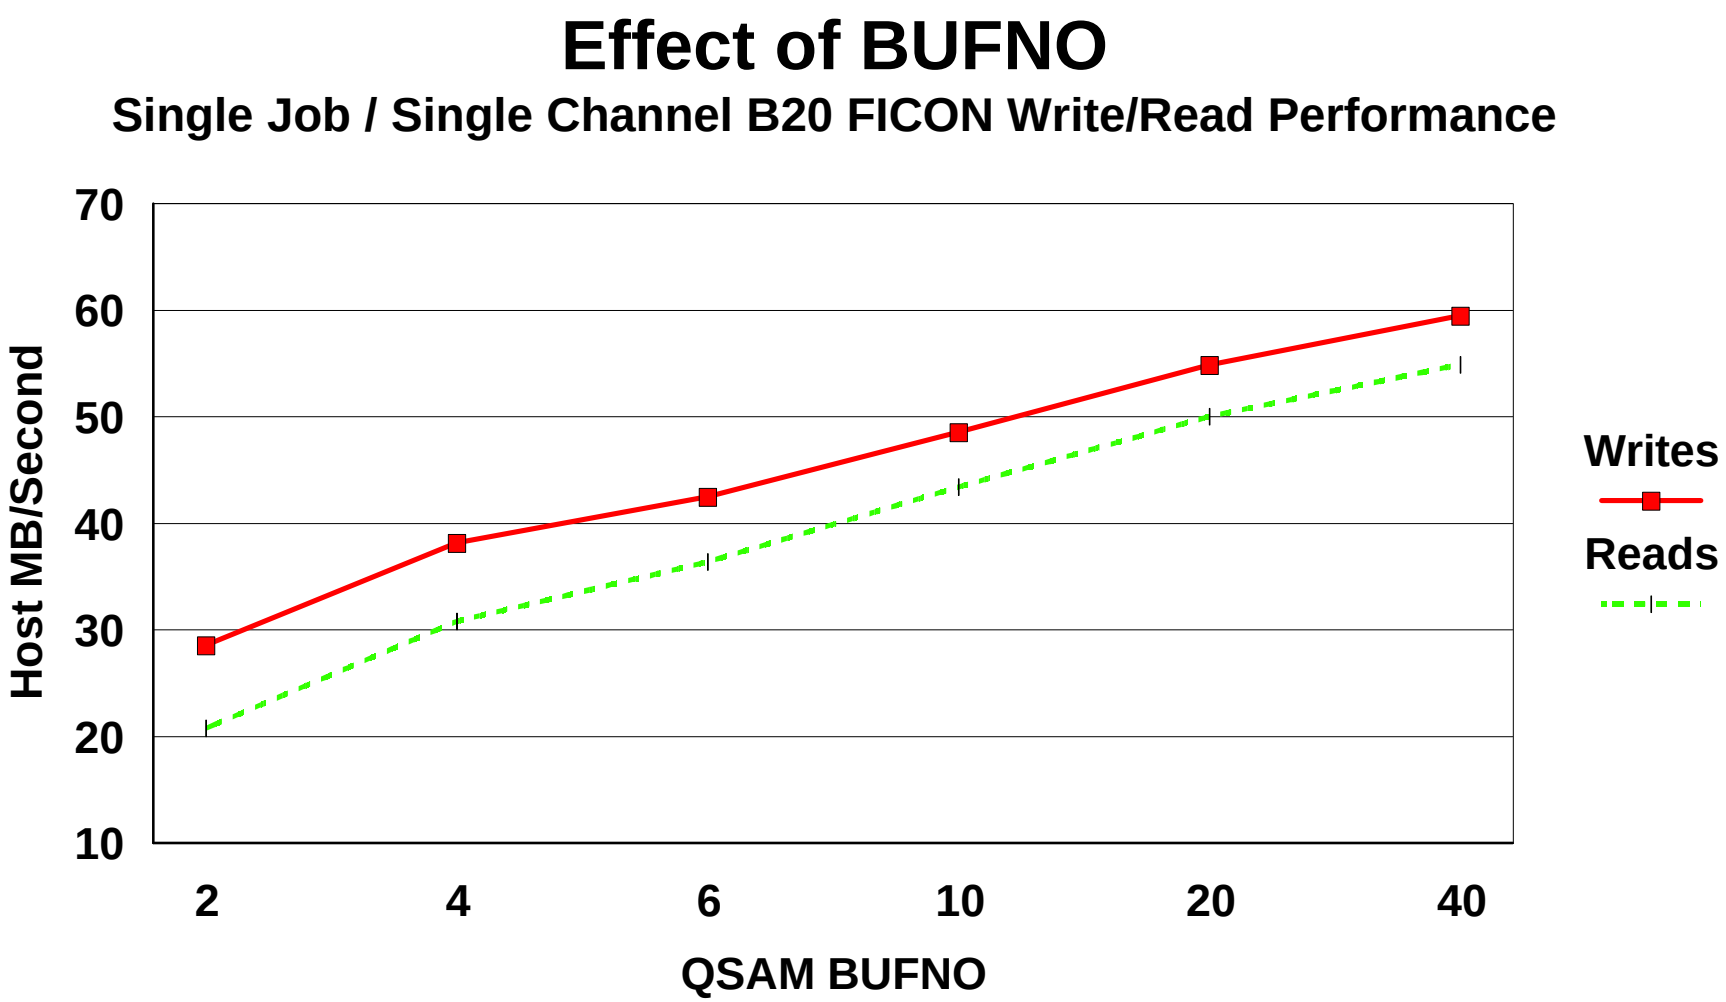

Chart 4: ESCON / FICON
Single-job Blocksize Performance

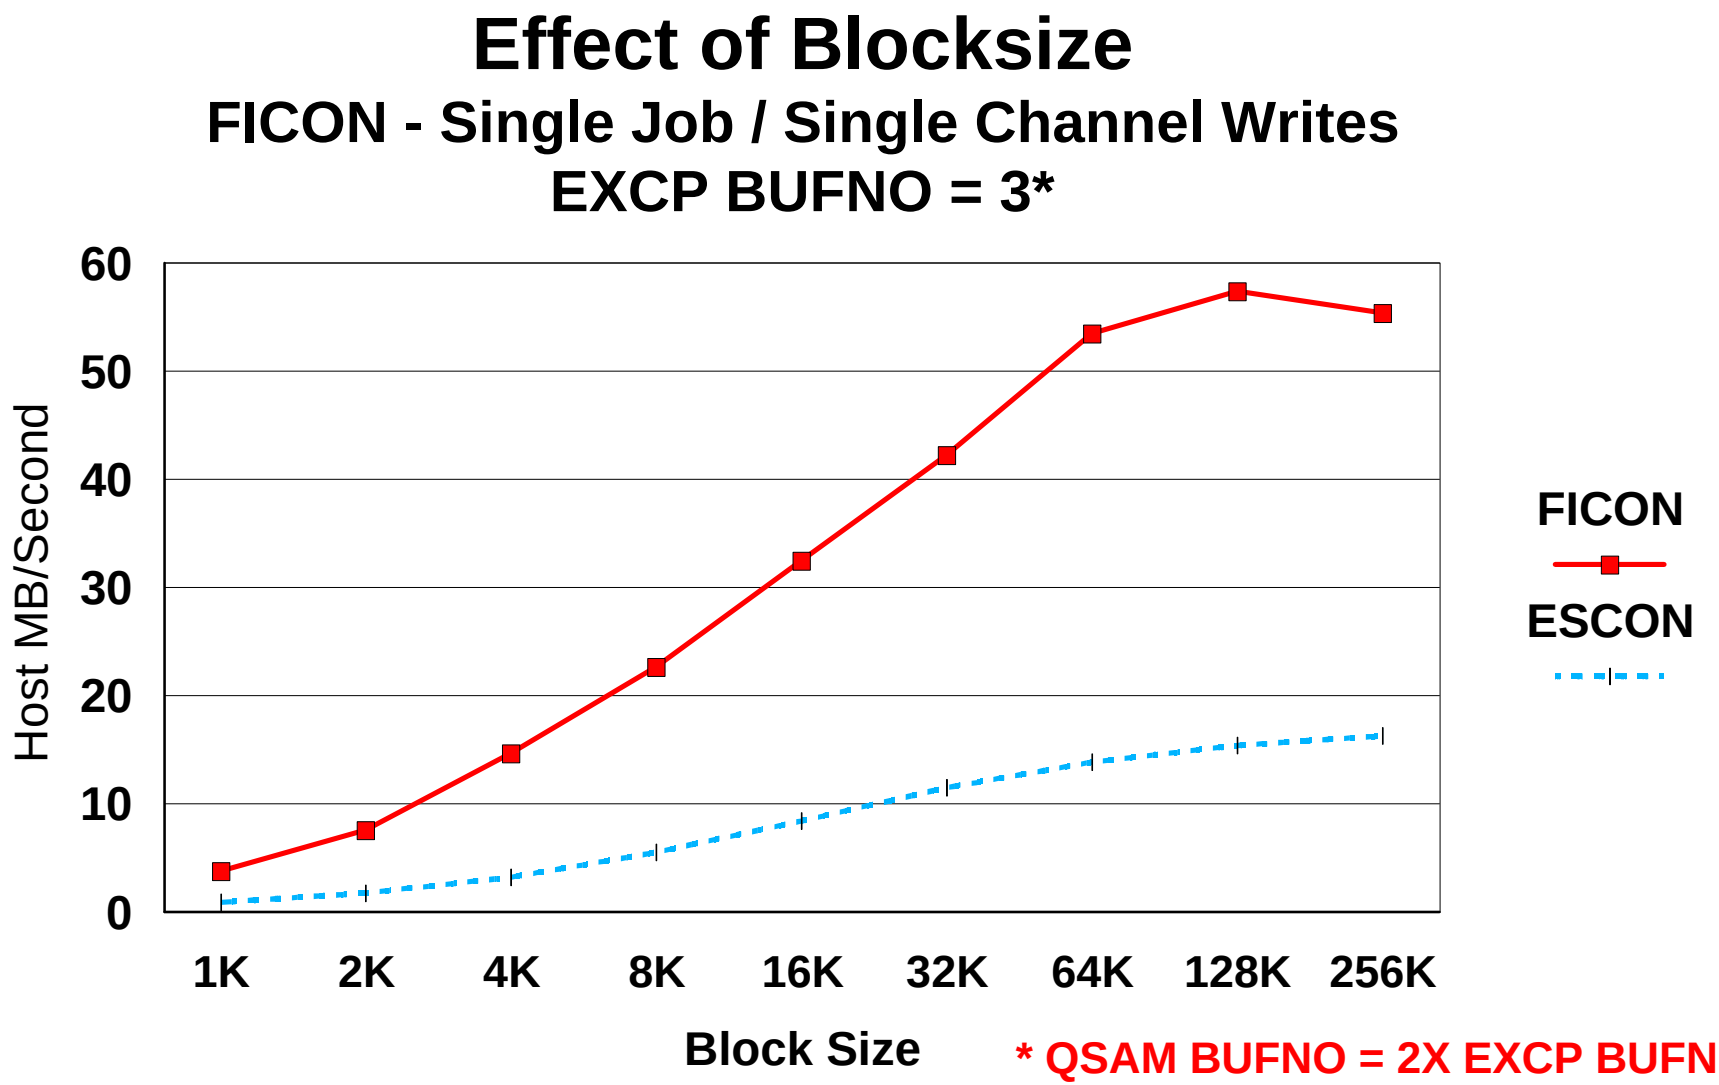

Chart 5: B10 ESCON / FICON
Multi-job Write and Read-Hit Performance

Model B10 ESCON vs. FICON Throughput
100% Host Write (Peak/Sustained) and 100% Read Hits

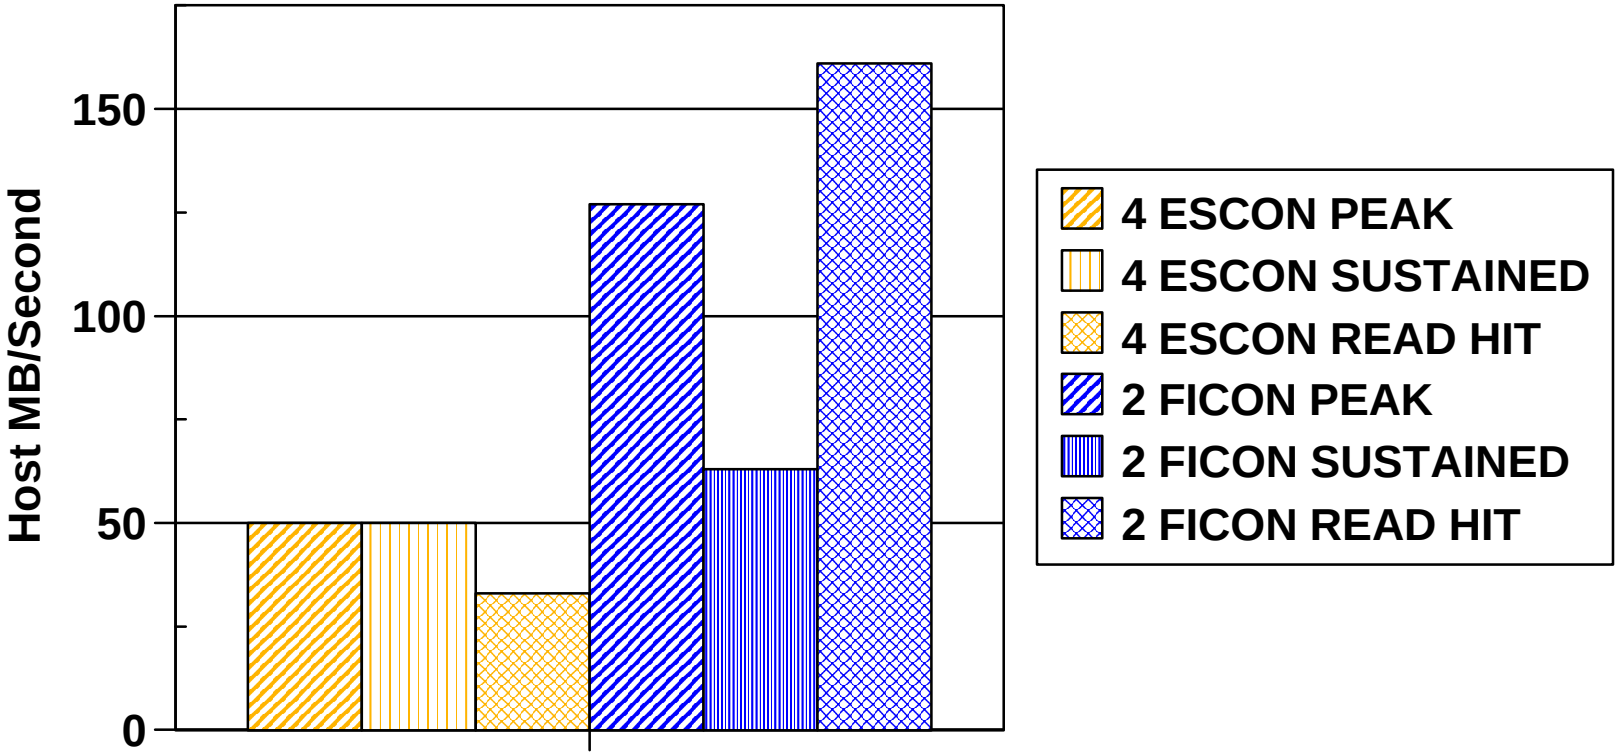

Chart 6: B20 ESCON / Base FICON
Multi-job Peak Write Performance

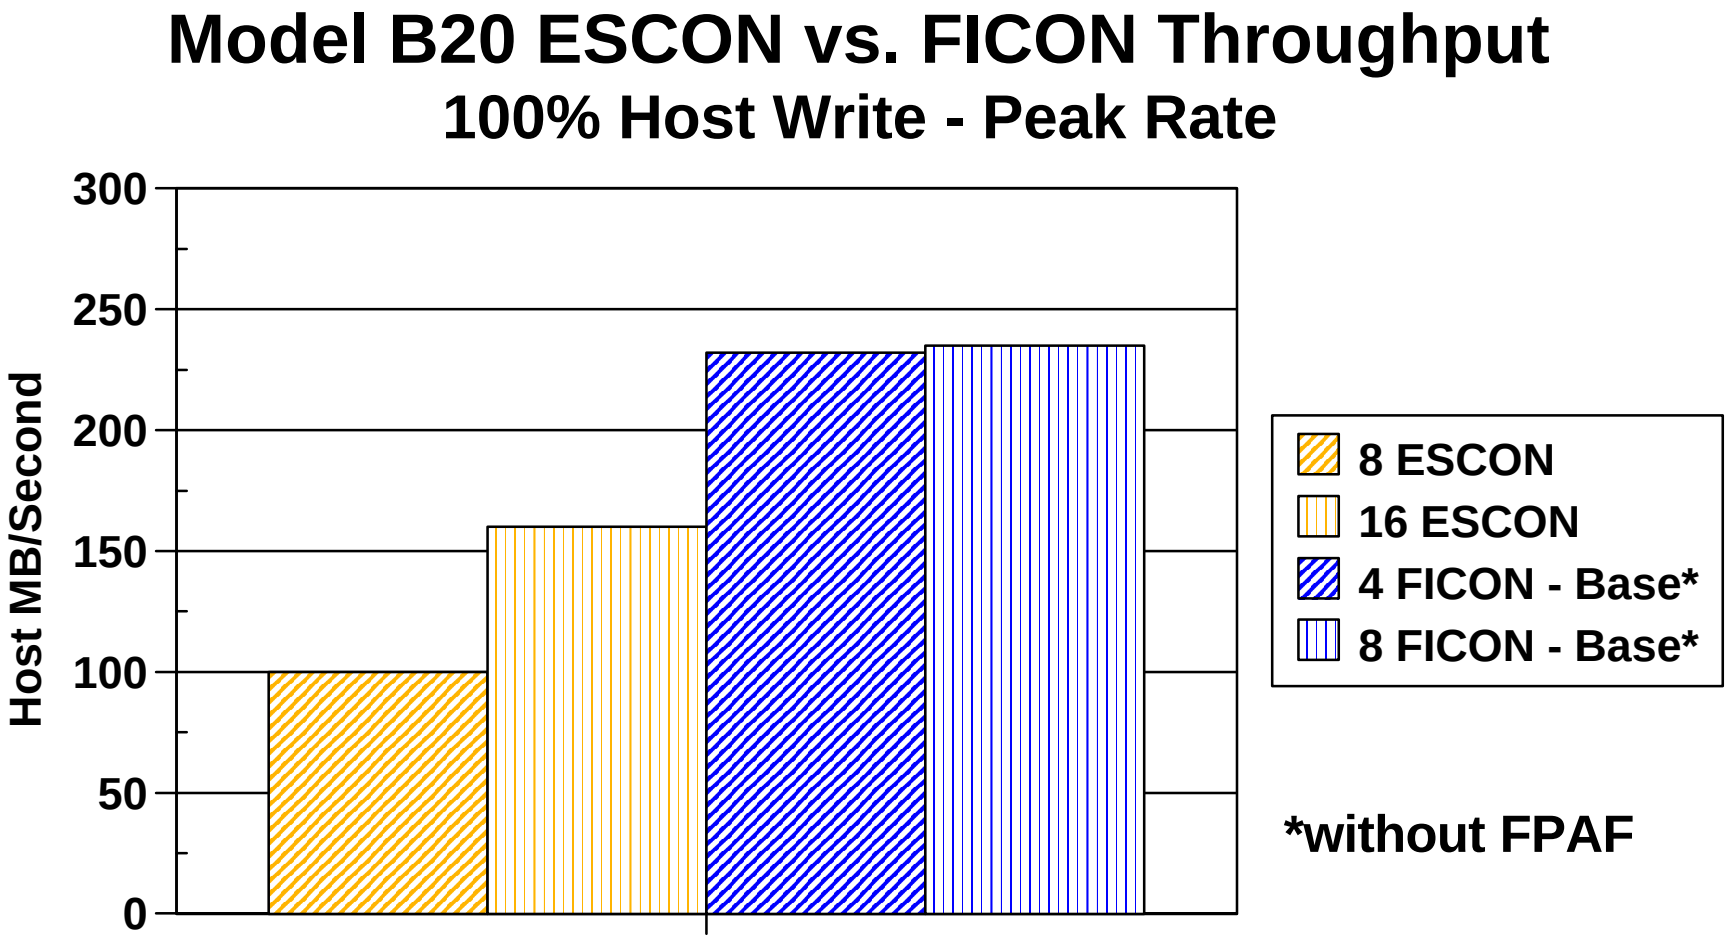

Chart 7: B20 ESCON / FICON PAF
Multi-job Peak Write Performance

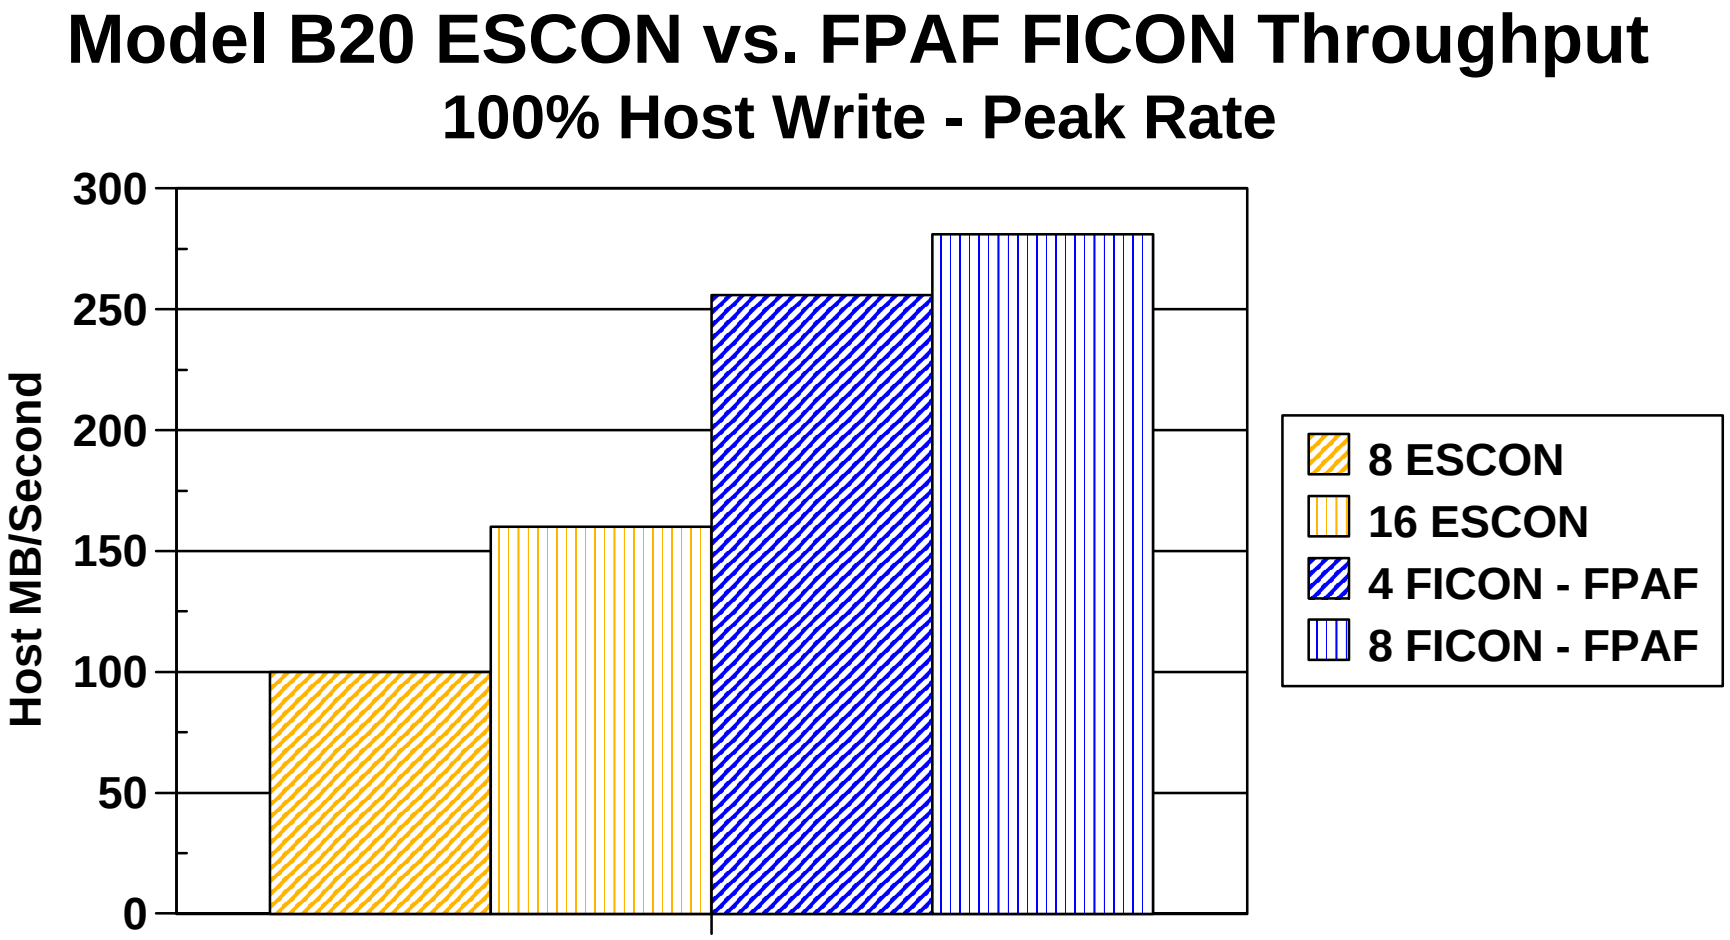

Chart 8: B20 Base / FICON PAF
Multi-job Peak Write Performance

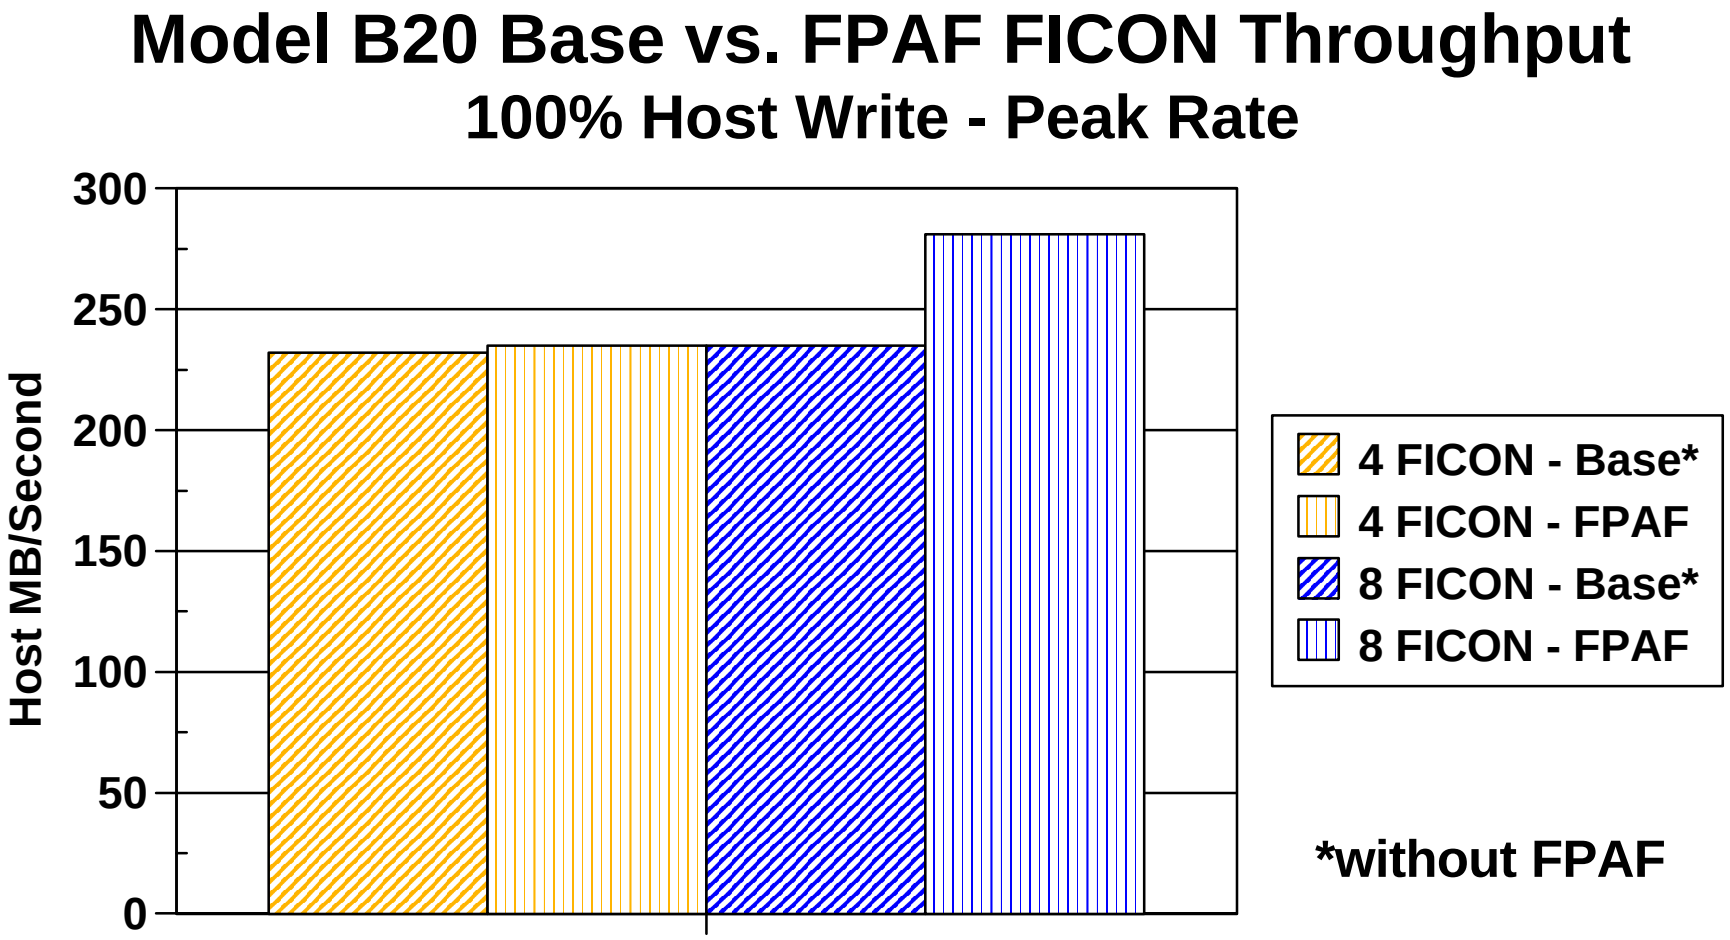

Chart 9: B20 ESCON / FICON
Multi-job Read-Hit Performance

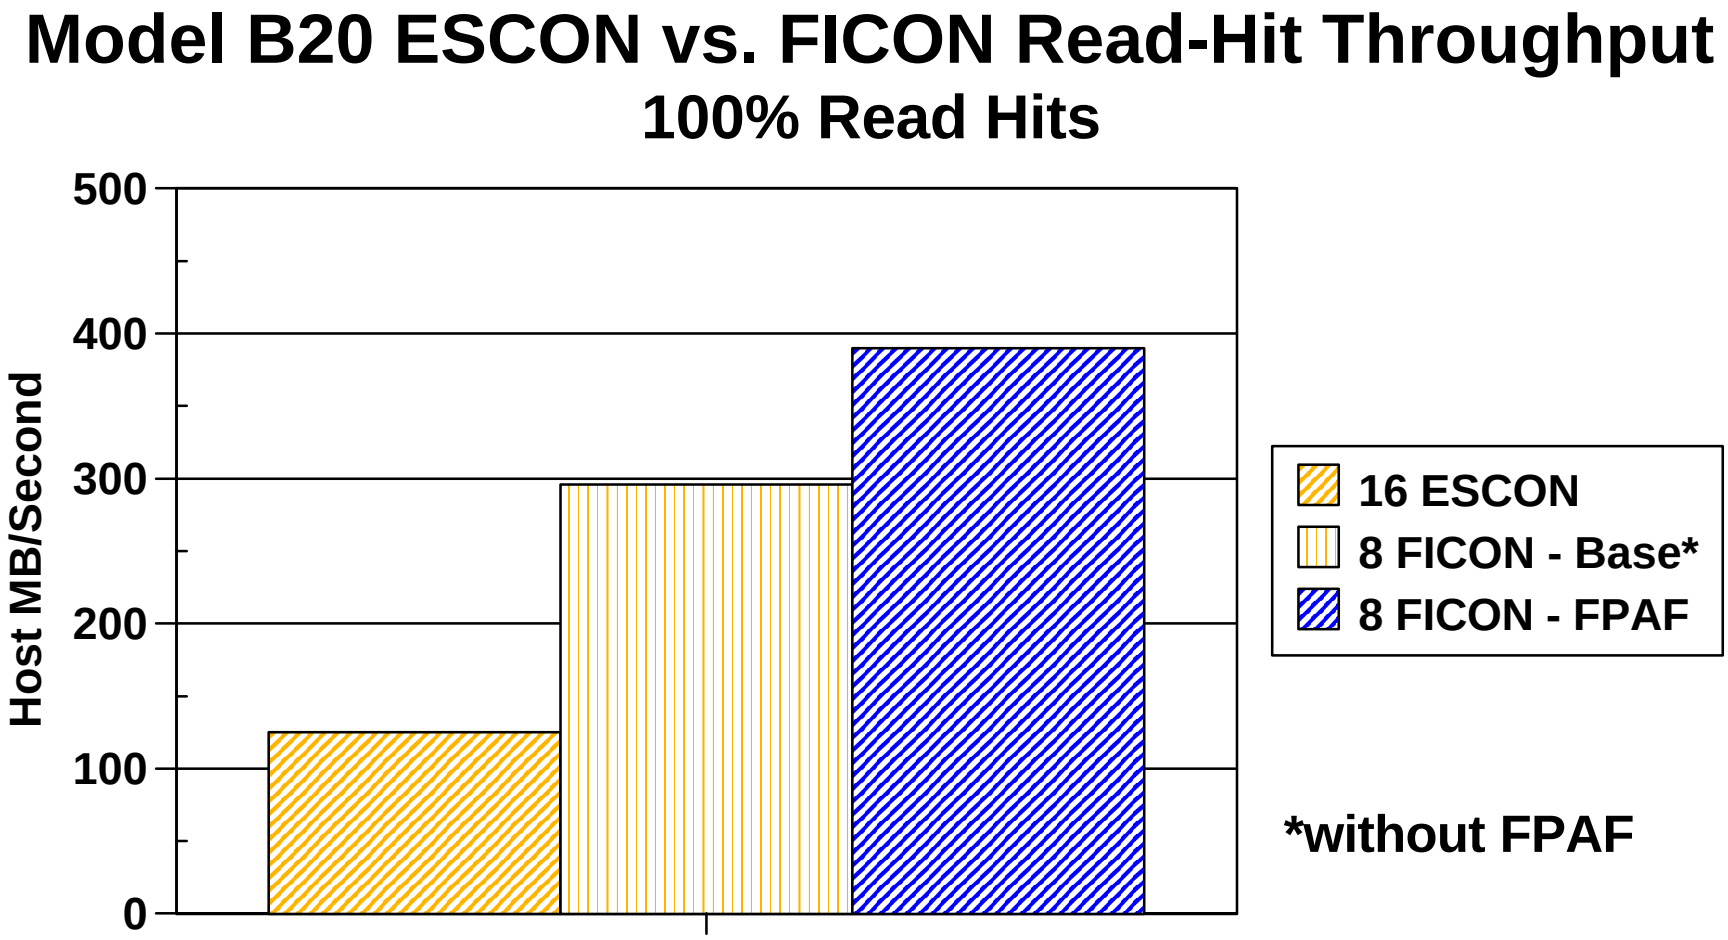

Chart 10: B20 ESCON / Base FICON
Multi-job Sustained Write Performance

Model B20 ESCON vs. Base* FICON Throughput
100% Host Write - Sustained Rate

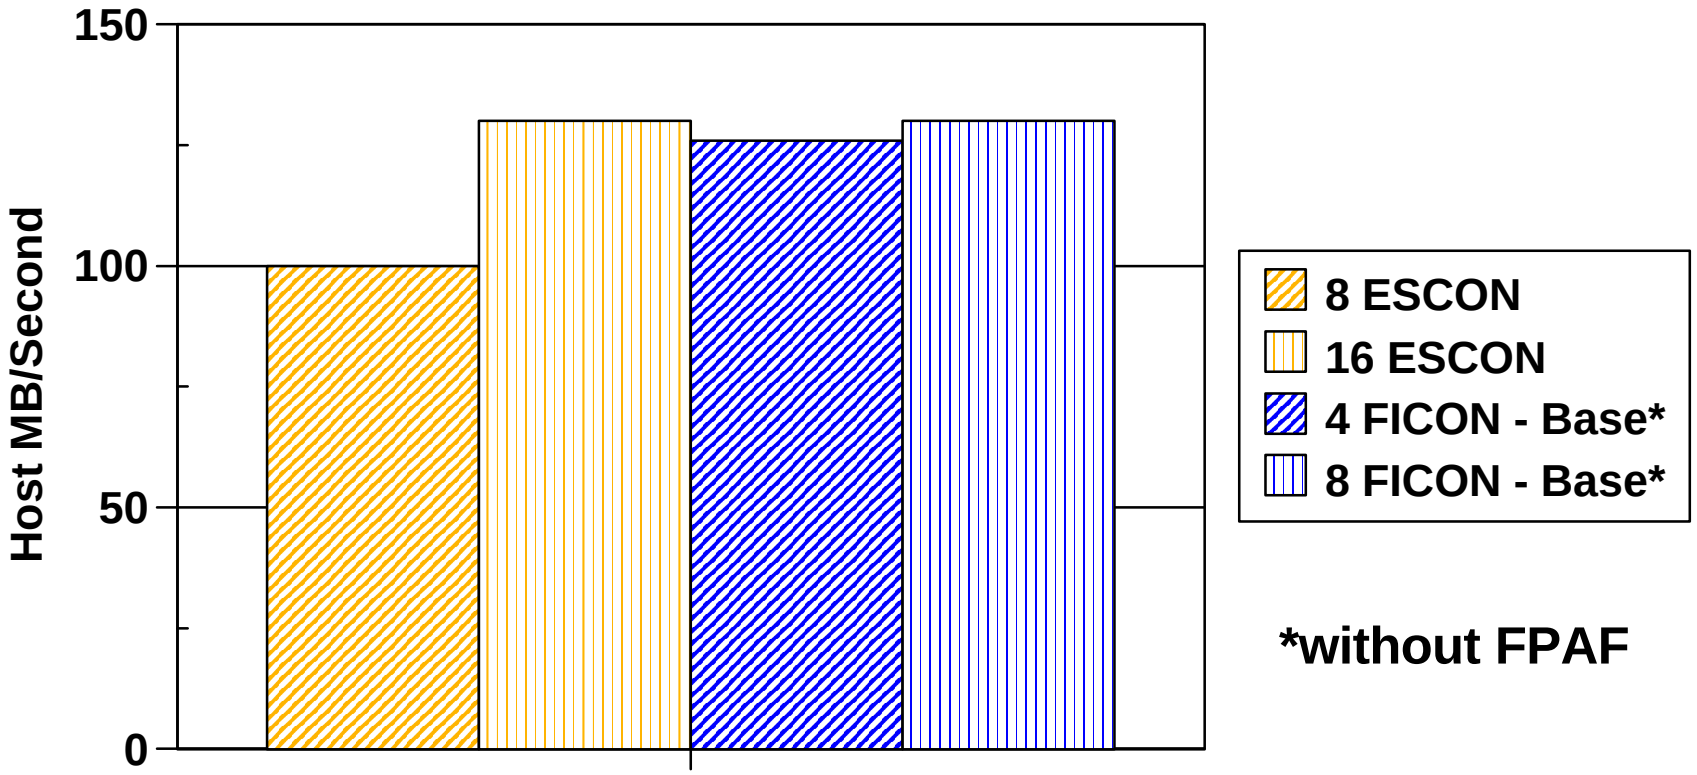

*without FPAF

Chart 11: B20 ESCON / FPAF FICON
Multi-job Sustained Write Performance

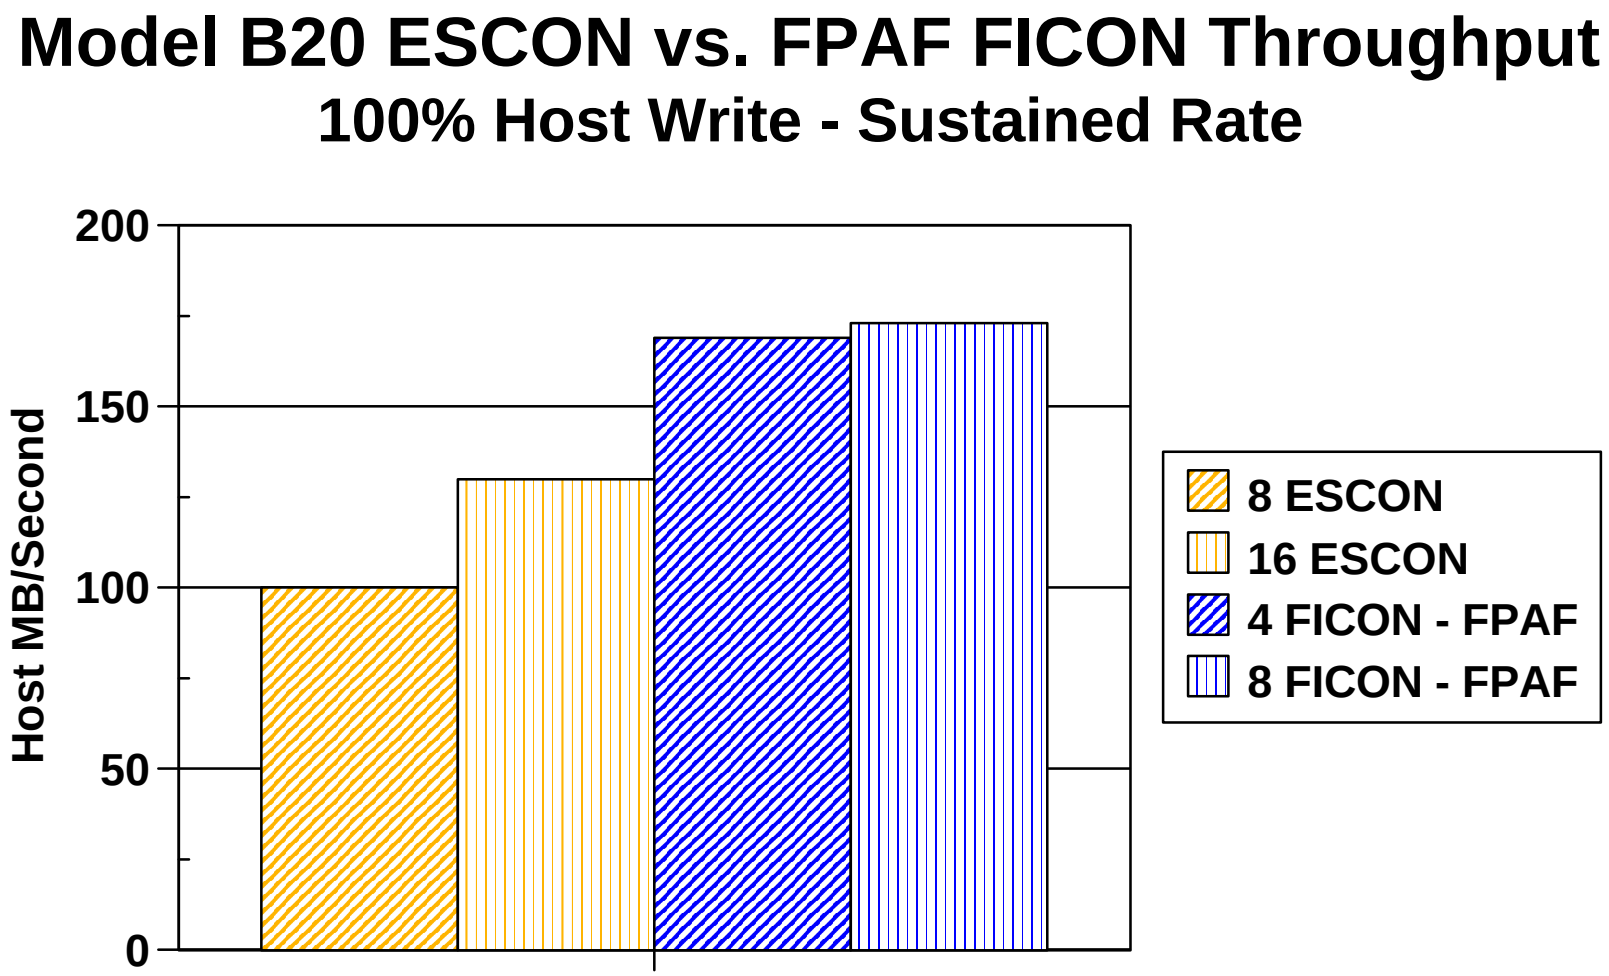

Chart 12: B20 ESCON / Base FICON
Multi-job Peak Performance vs. # Physical Drives

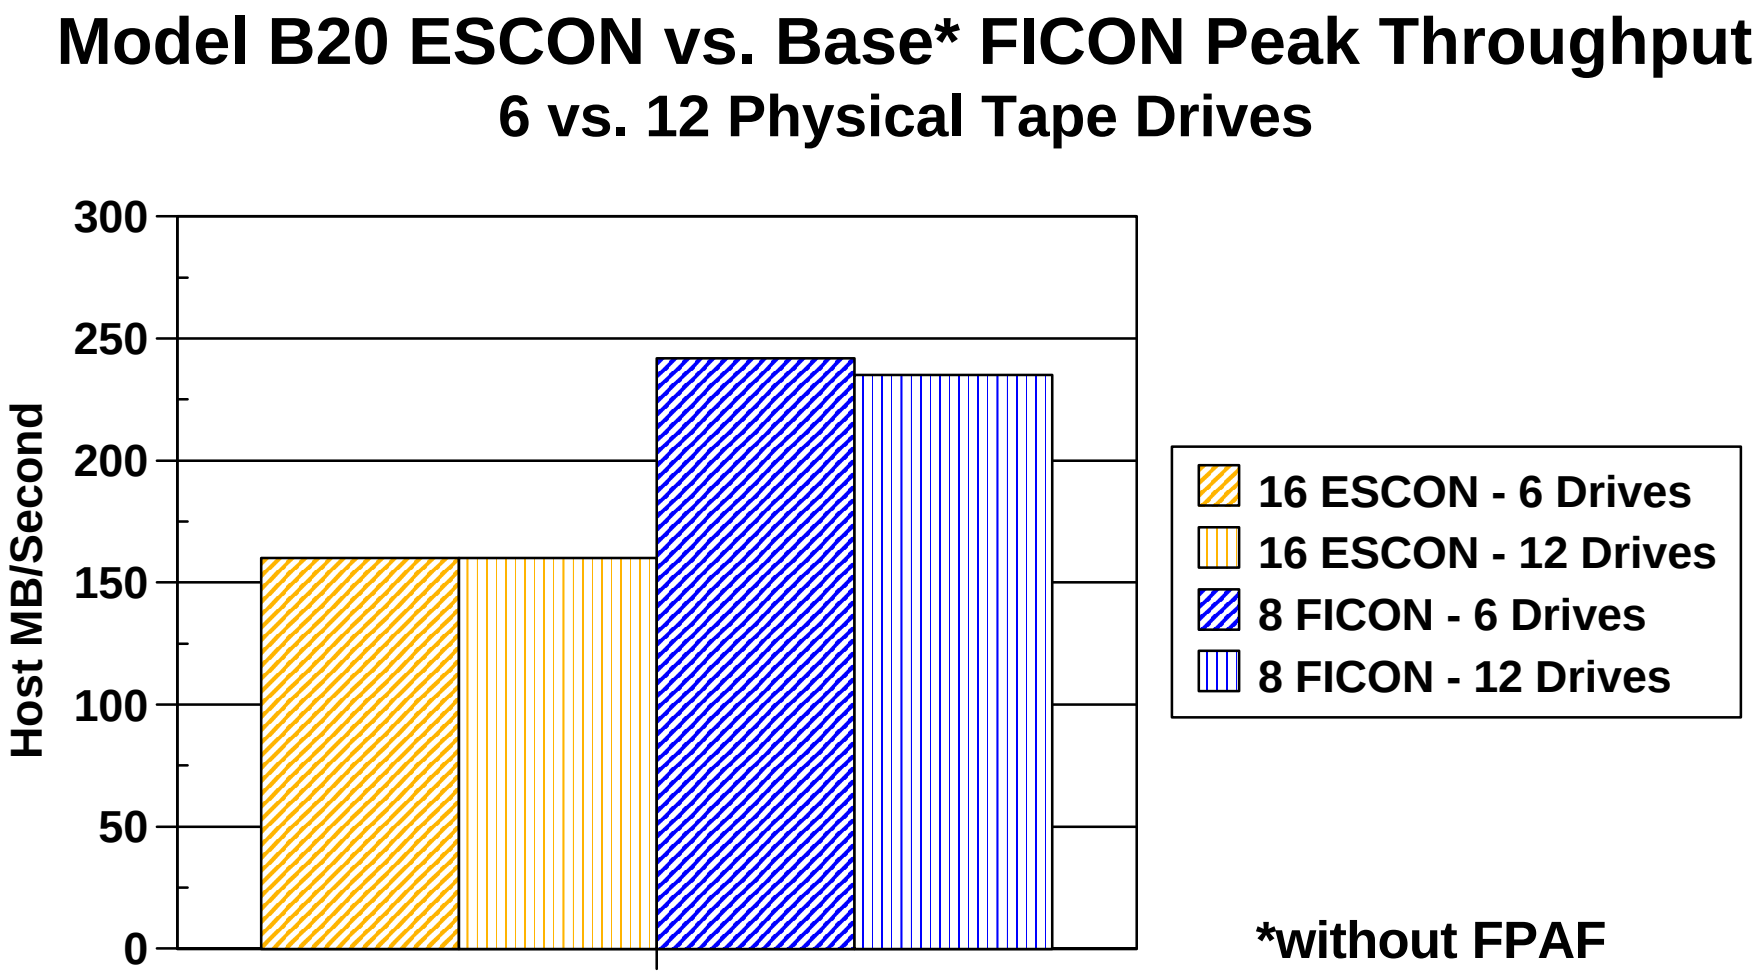

Chart 13: B20 ESCON / Base FICON
Multi-job Sustained Performance vs. # Physical Drives

Model B20 ESCON vs. Base* FICON Sustained Throughput
6 vs. 12 Physical Tape Drives

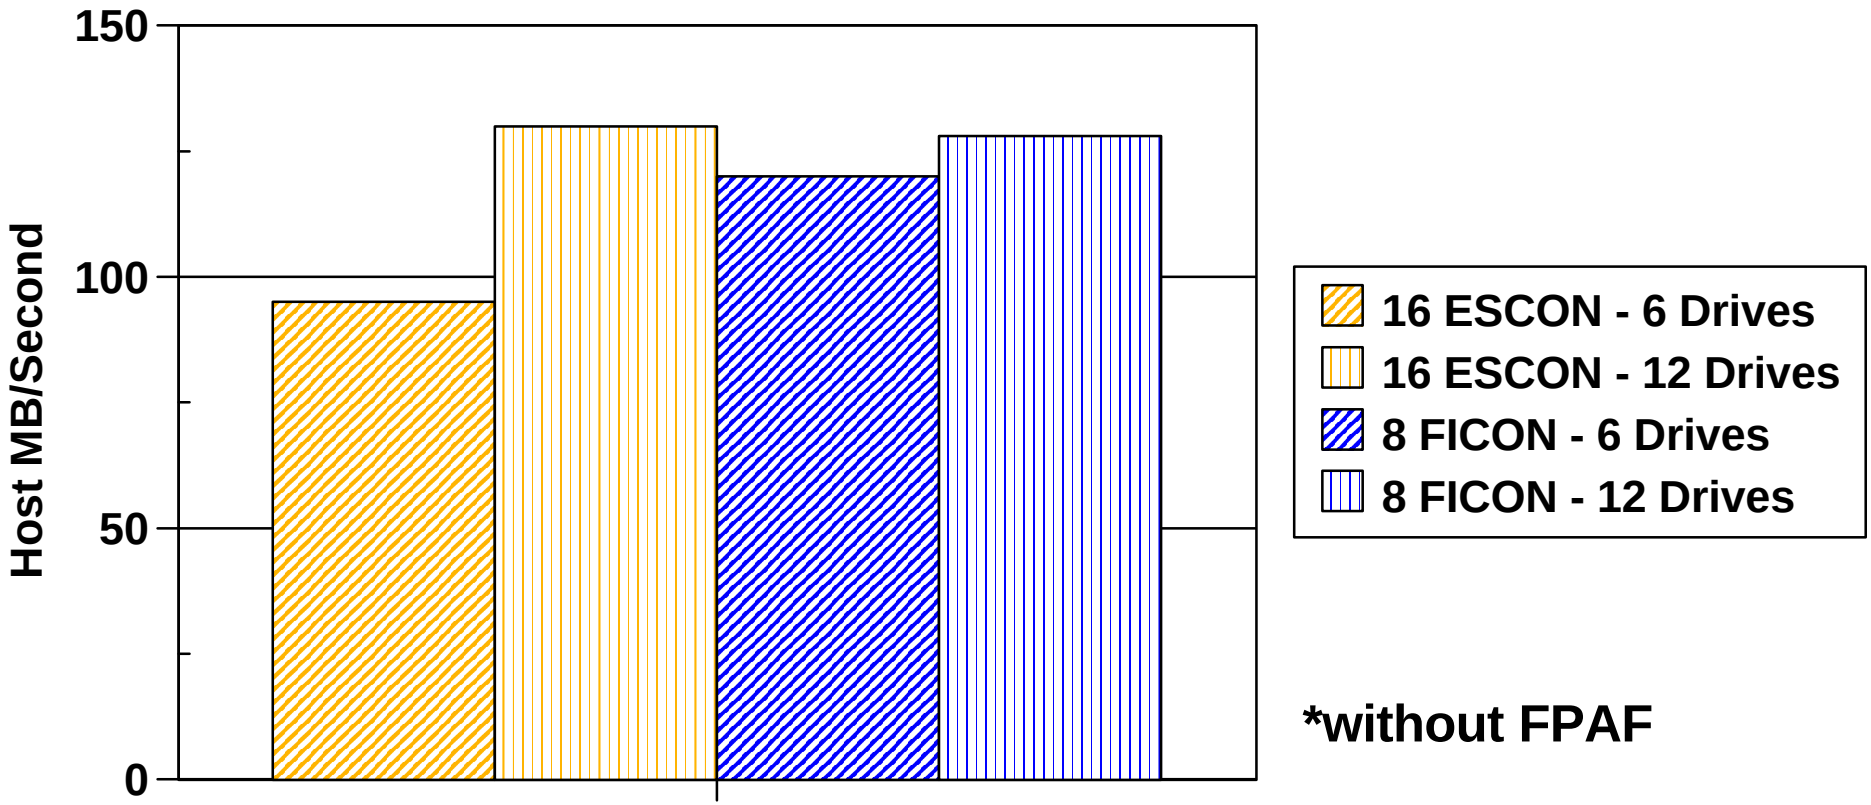

*without FPAF

Chart 14: B10 ESCON / Base FICON
Multi-job Peak Write Performance - BUFNO=5

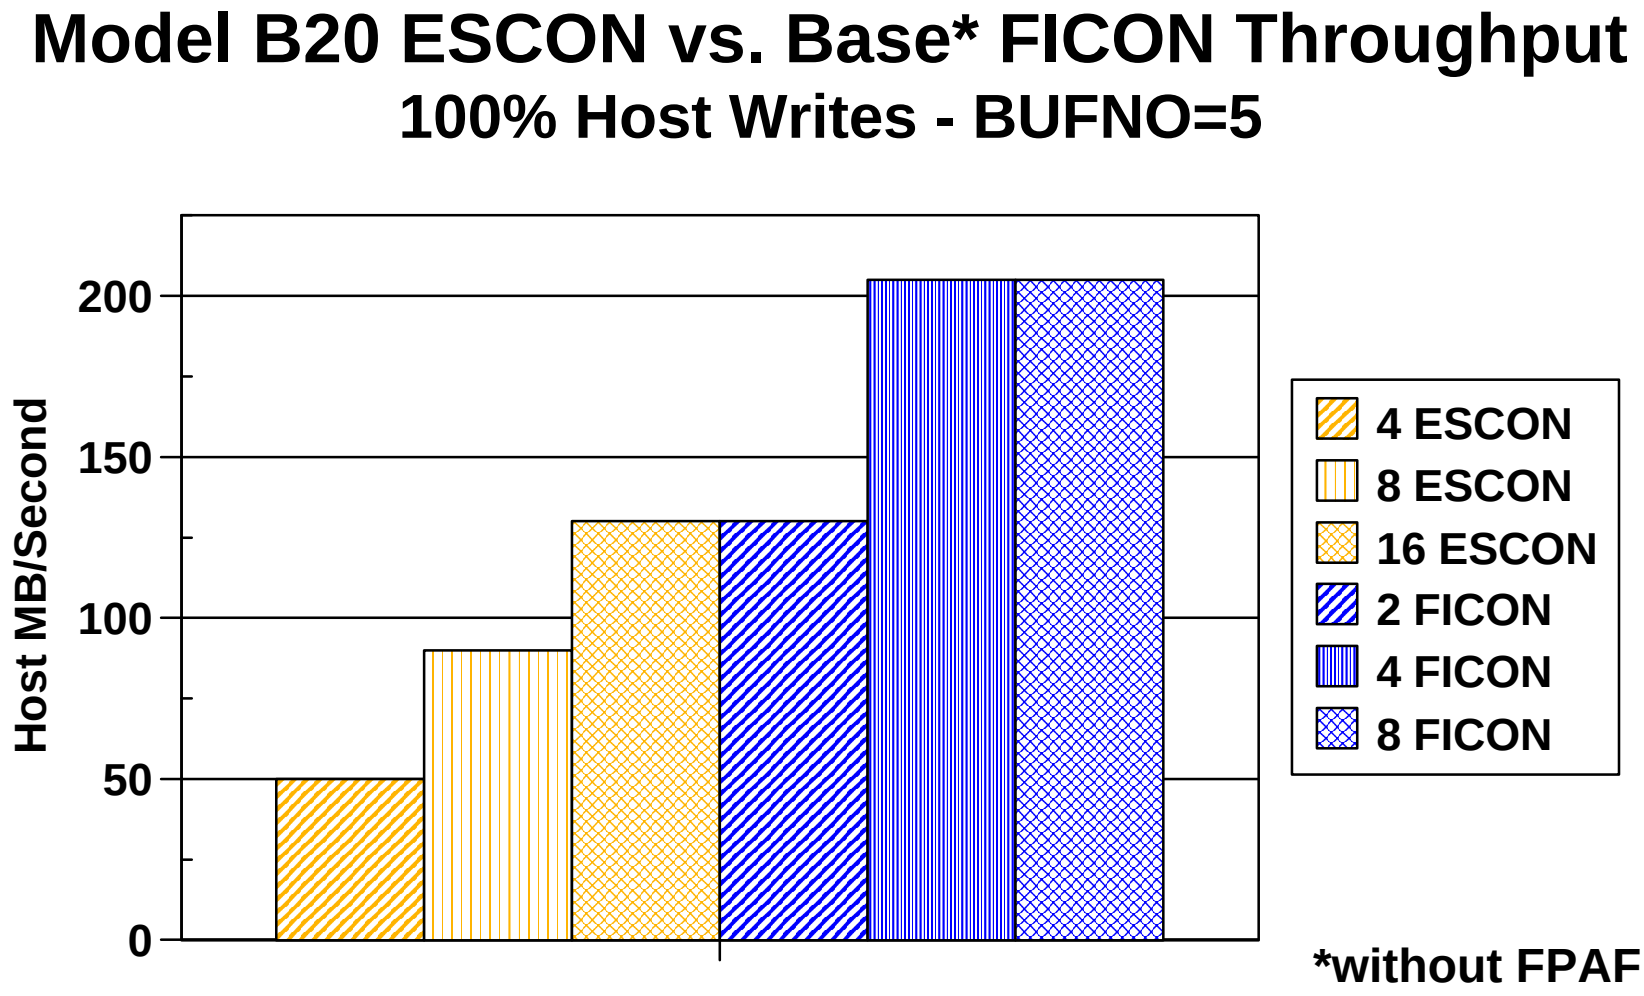

Chart 15: B20 FICON
Multi-job Peak Write Duration

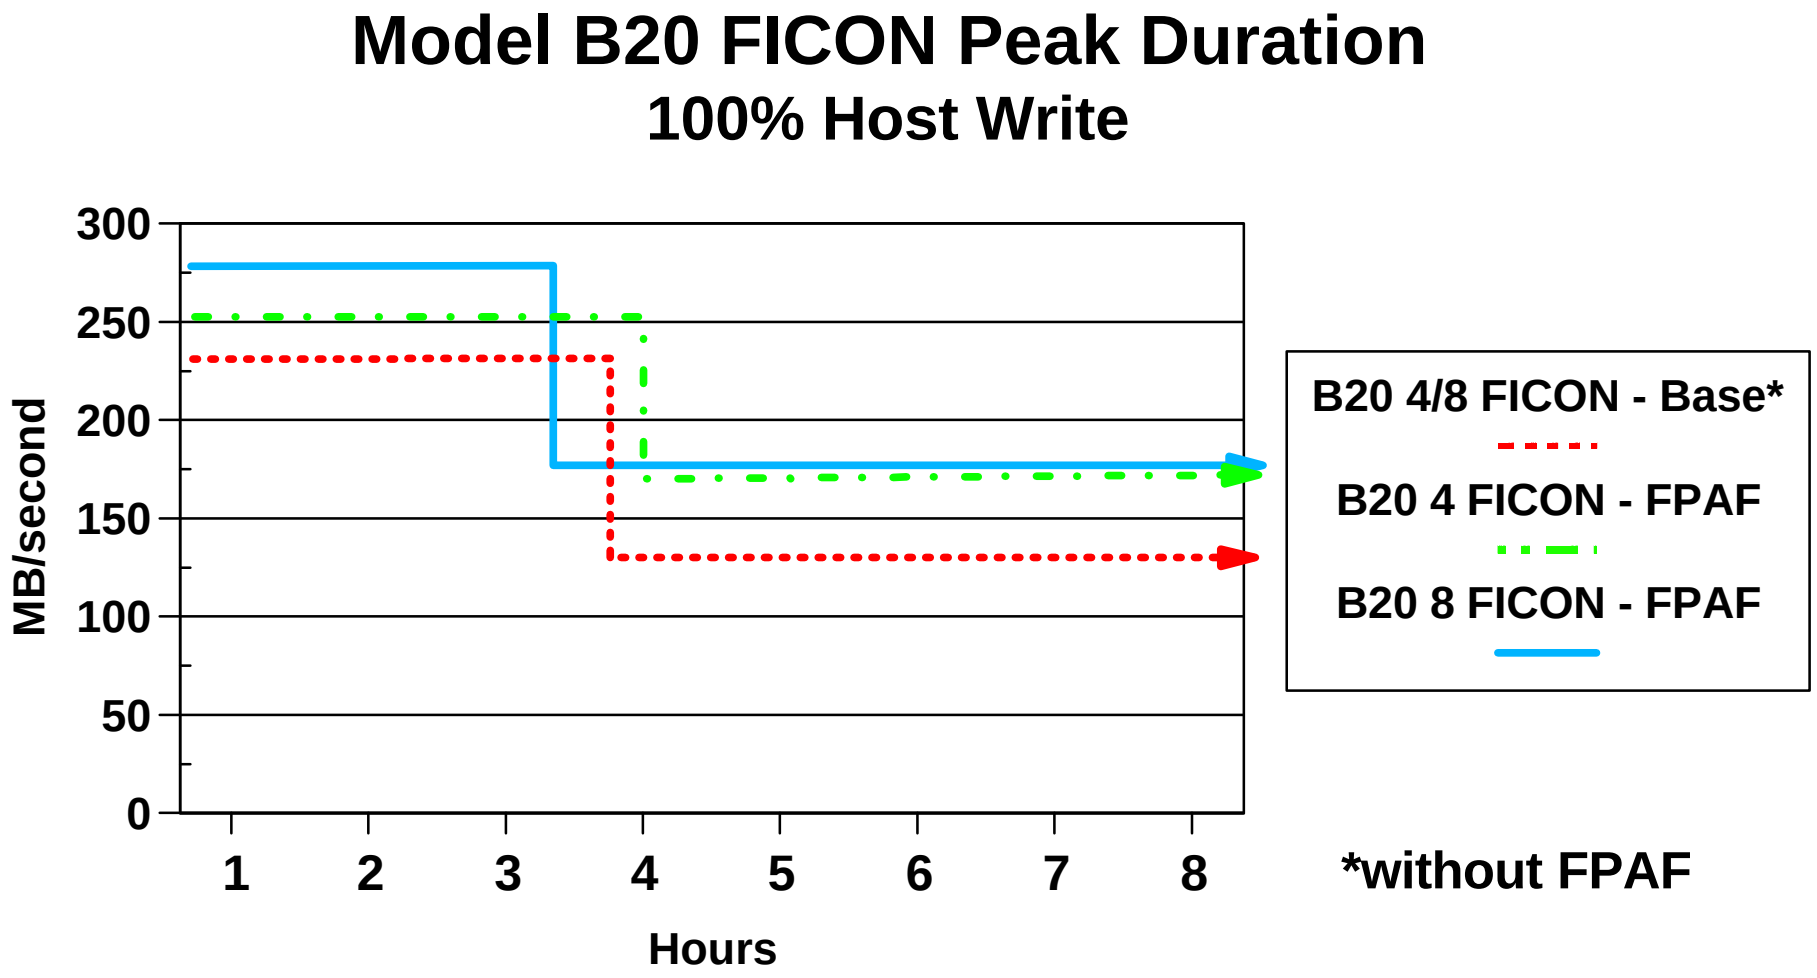

Chart 16: B10 FICON
Multi-job Write Performance vs. Data Compression

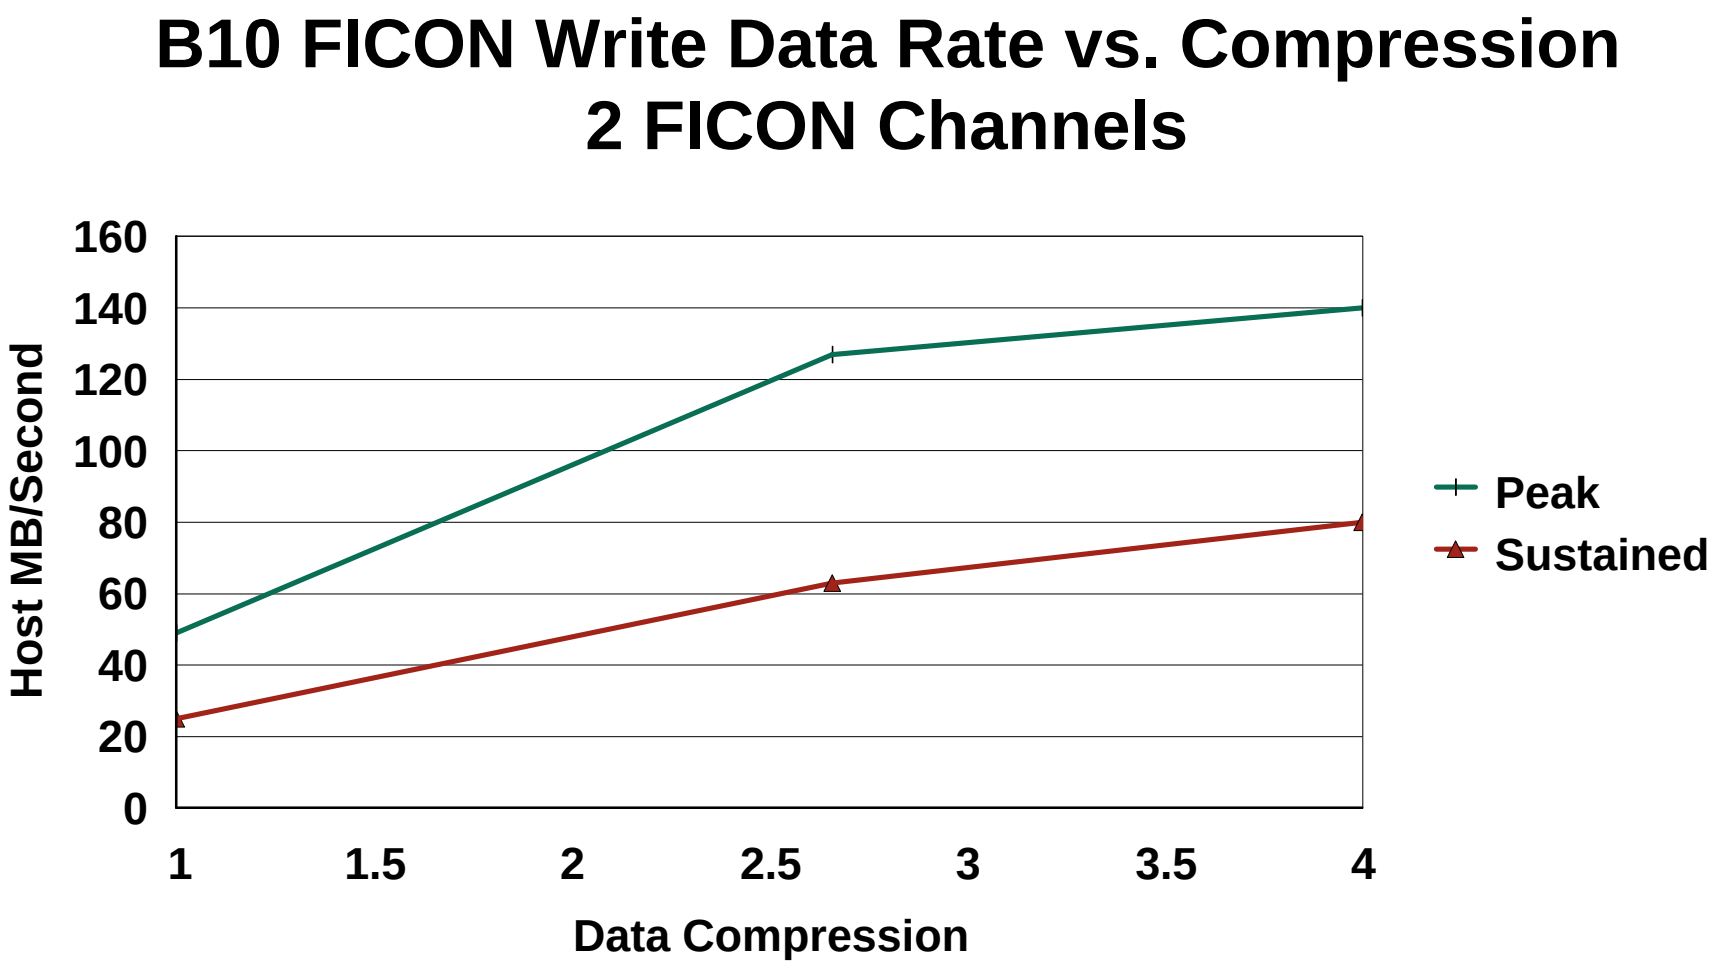

Chart 17: B20 Base FICON
Multi-job Write Performance vs. Data Compression

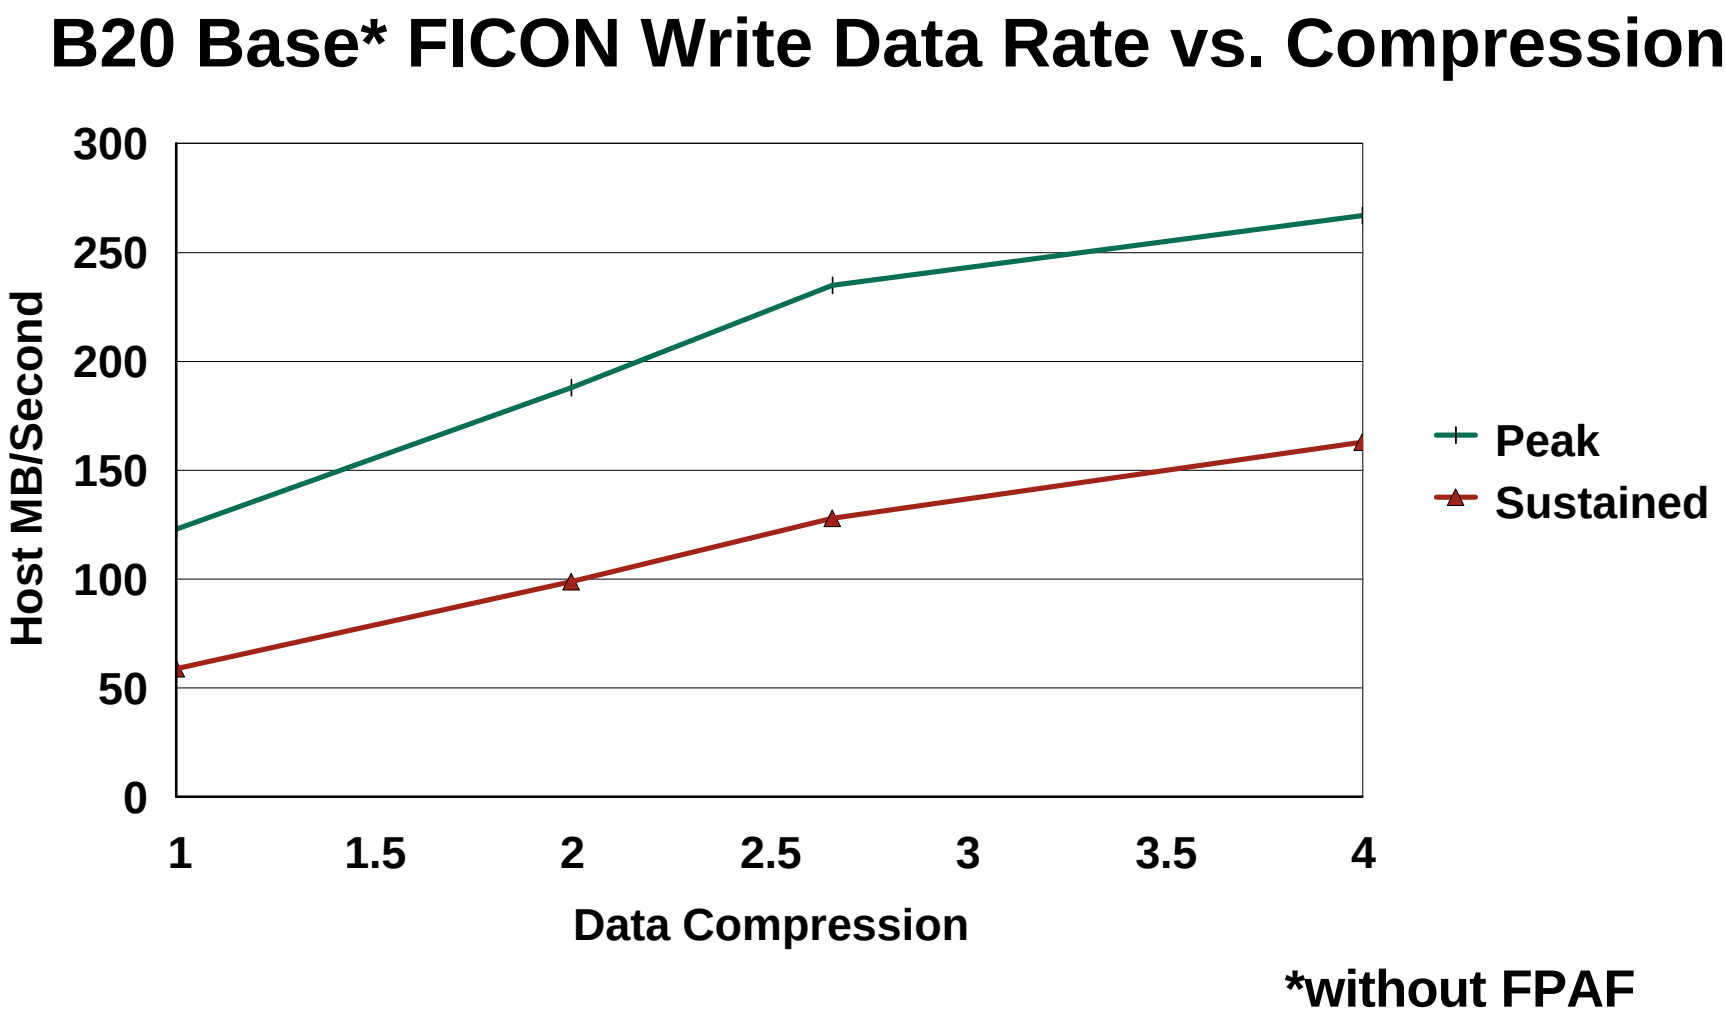

*without FPAF

Chart 18: B20 Base FICON
Multi-job Write Performance vs. # Virtual Drives

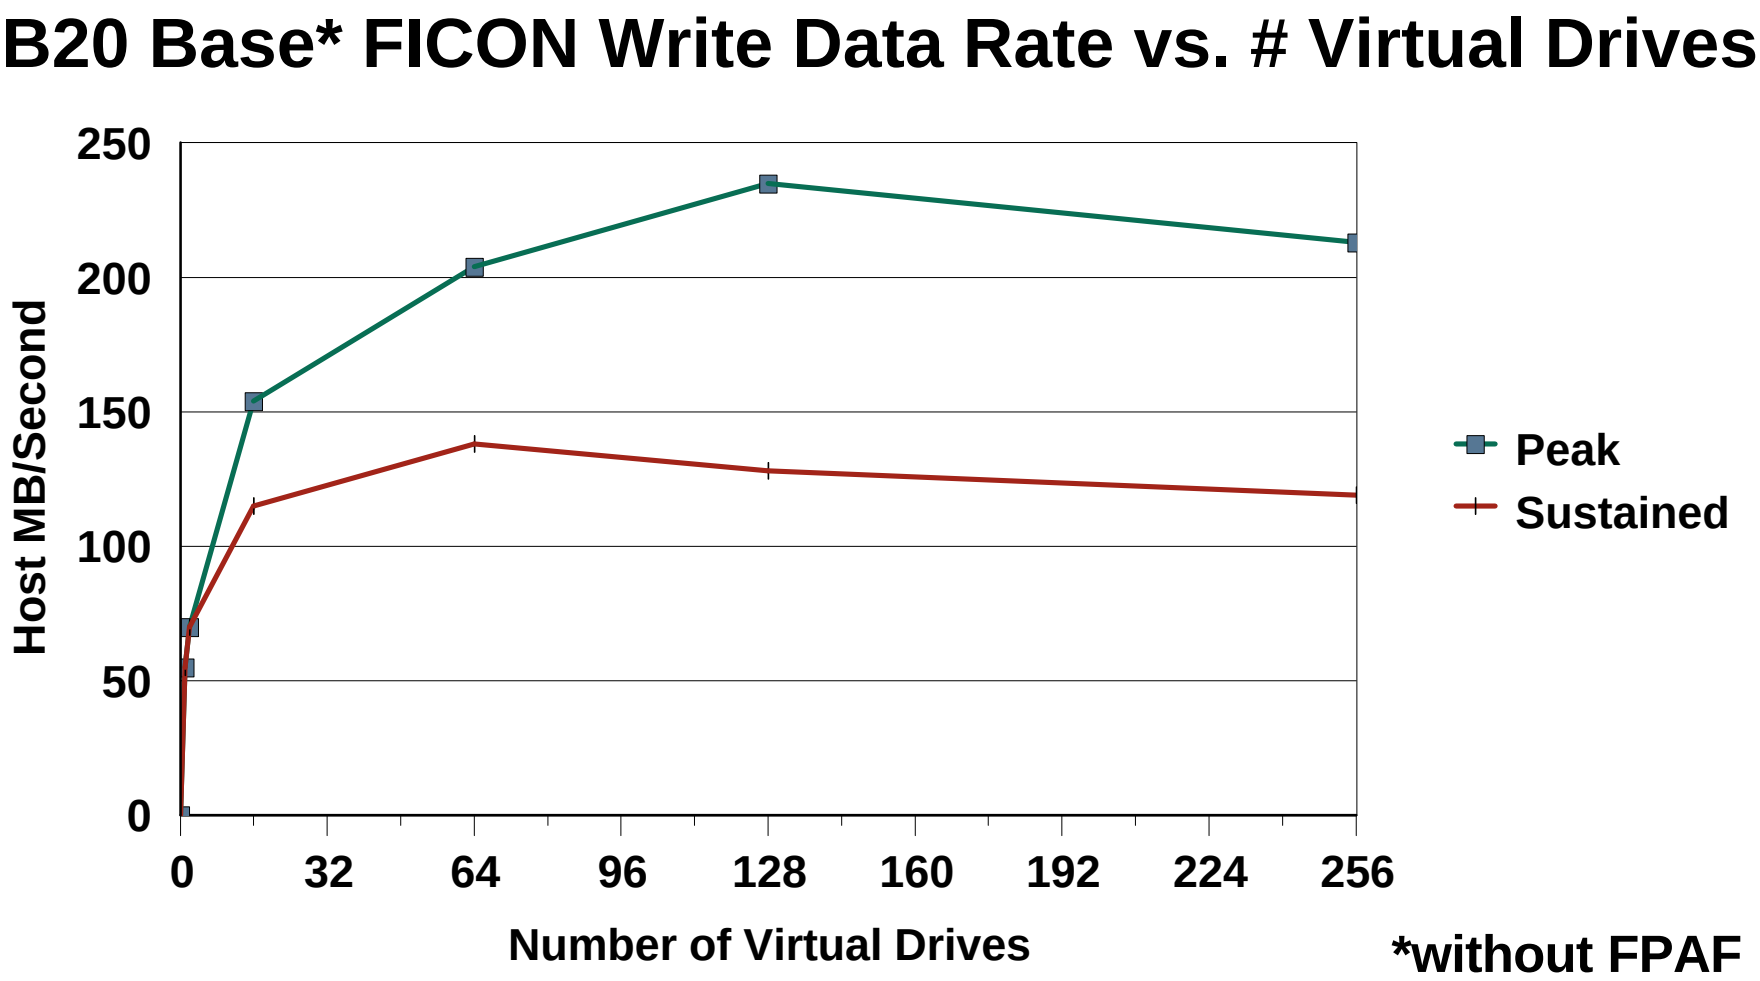

Chart 19: B20 Base FICON
Multi-job Write Performance vs. Blocksize

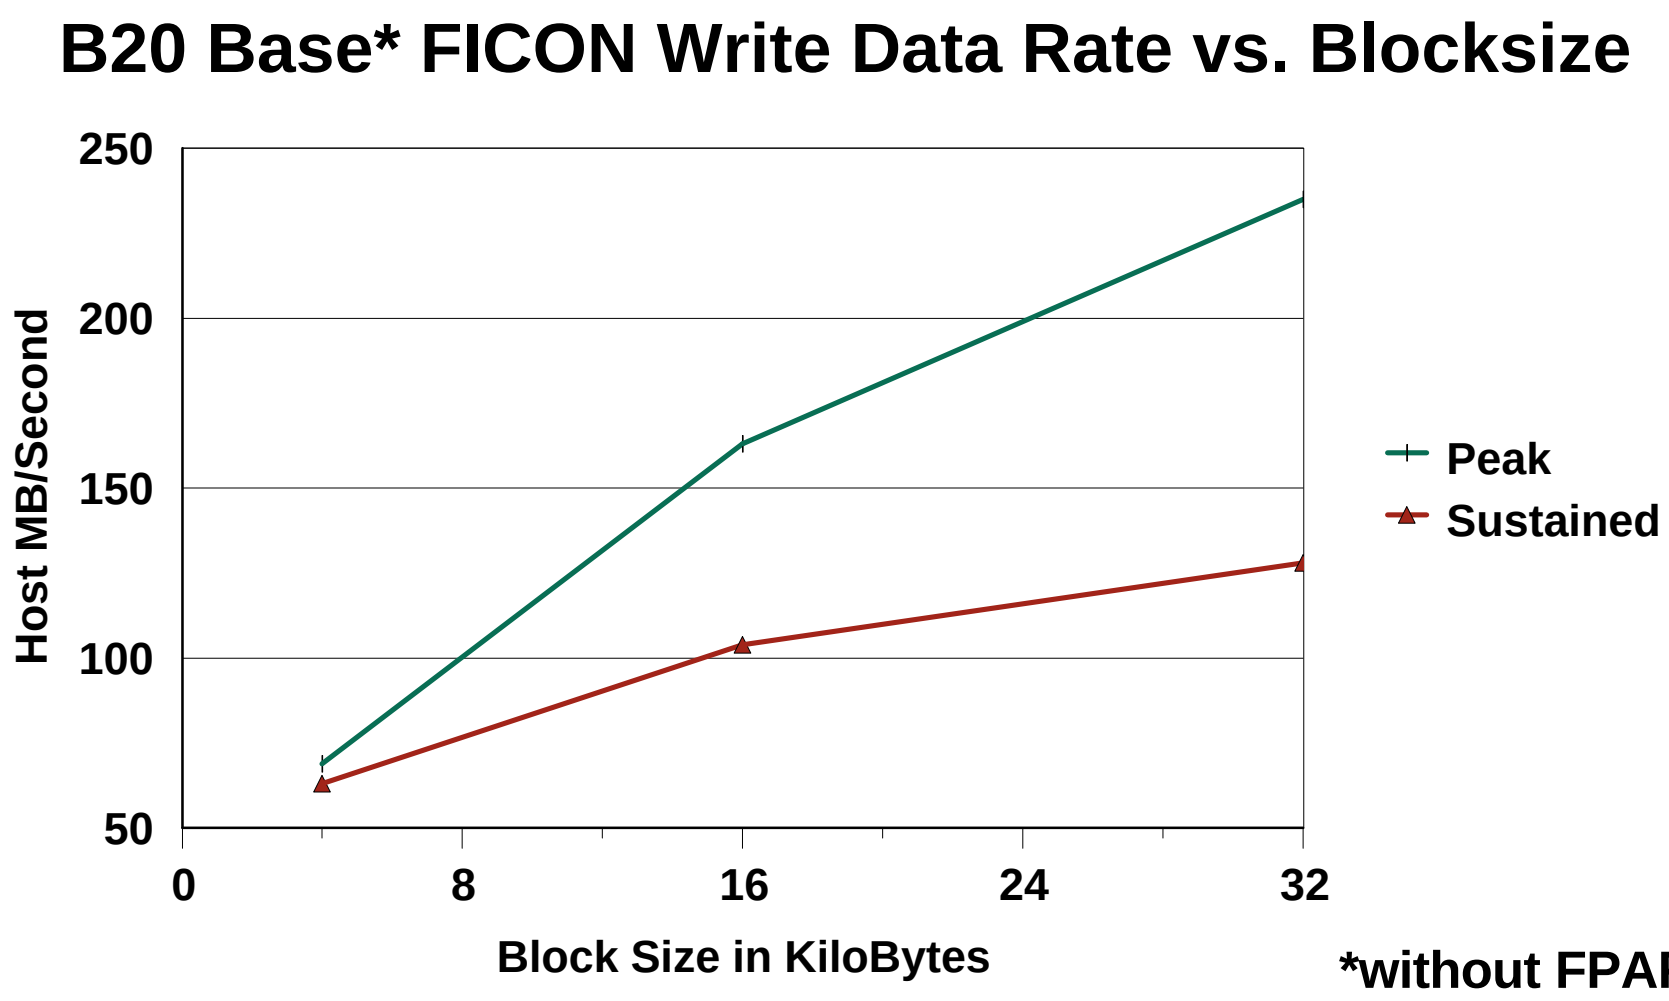

Chart 20: B20 Base FICON
Multi-job Write Performance vs. Volume Size

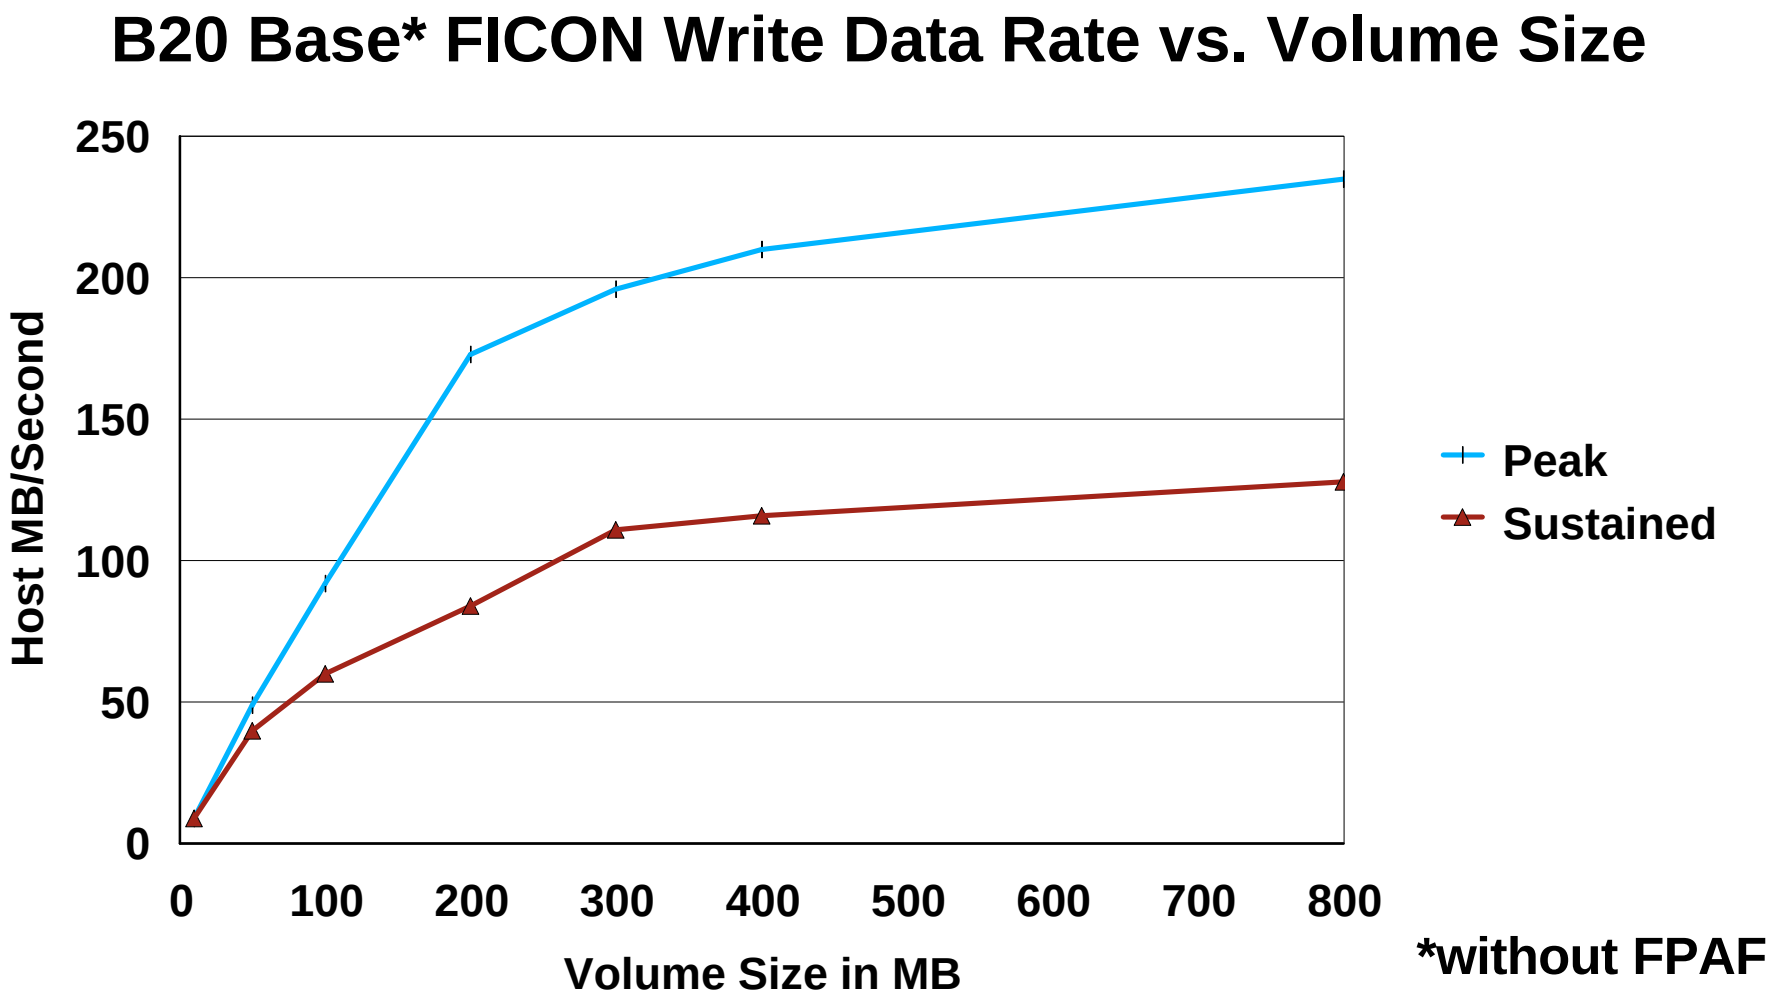

Chart 21: B20 Base FICON
Multi-job Read Performance vs. Data Compression

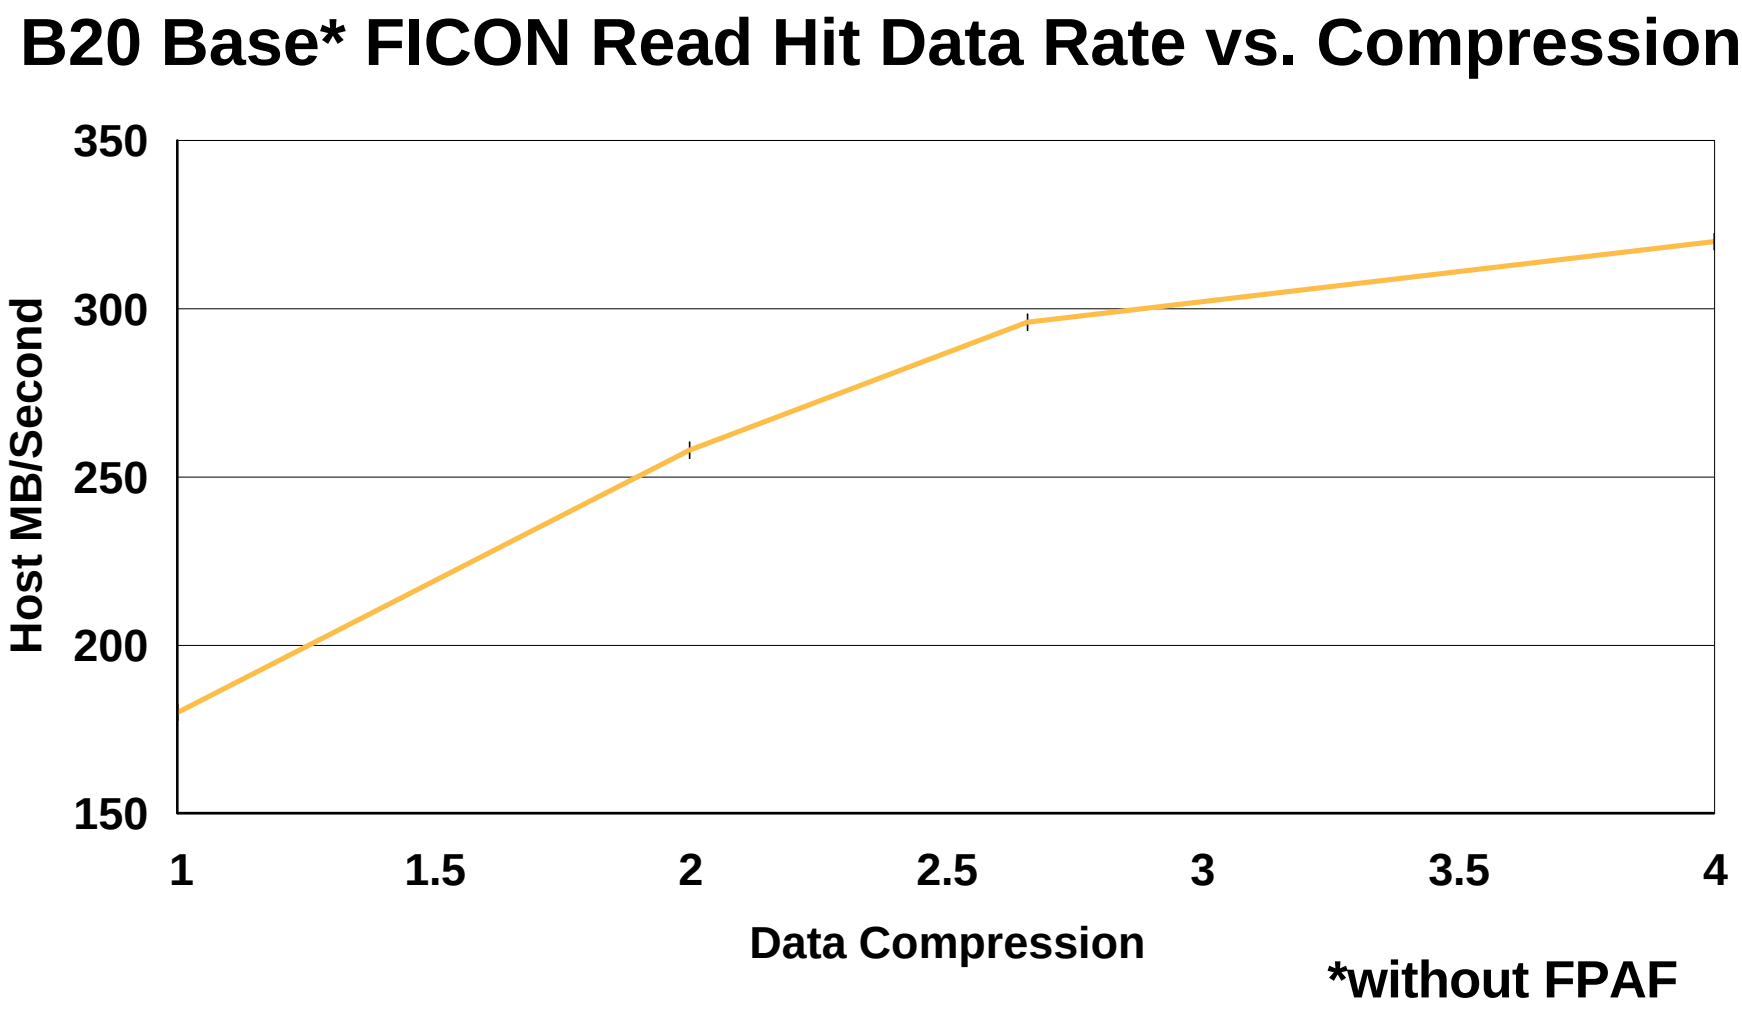

VTS FICON Miscellany

- **Single job B20 performance should be similar to single job B10 performance.**
- **Benefits of large BUFNO value are greater for FICON than for ESCON.**
- **For more than two FICON attachments, split CHPIDs over two or more processor IOPs (at least for a G6).**
- **Initial measurements indicate that B10 4-FICON channel performance should be 15% - 20% higher than 2-FICON channel performance at 2.66:1 compression.**
- **B20 measurements in this document were made with four cache drawers (1.7 TB). Limited FICON Base measurements with a two-drawer cache show no significant performance difference with data compression of 2.66:1. However, it is expected that the high data rates shown for FICON PAF will only be achieved with four cache drawers.**
- **All local measurements were made using a G6 host processor. Initial measurements indicate that a Freeway processor can drive a VTS FICON channel at a 20% to 30% higher data rate than can a G6 processor. Such an increase is not expected to affect most multi-job data rates shown in this document.**

VTS Recall Considerations

Recall rates have always been a challenge to characterize, because of the potentially wide variations in performance that can occur due to both the placement of the recalled logical volumes on the stacked volumes and on the order in which the logical volumes are recalled.

To minimize the effects of volume pooling on sustained write performance, the VTS code which supports APM implements a new algorithm that blocks several volumes together on the same stacked volume for premigration. This new algorithm is used whether or not APM is enabled or volume pooling is implemented. As a result, some customers may see a somewhat different level of recall performance with this VTS release. An example is given on the next page of the difference this new algorithm might make in logical volume distribution over physical volumes

Basic recall throughput performance should not change for FICON vs. ESCON channels, due to the fact that the physical backend drives remain the primary limiting performance factor.

Logical Volume Distribution over Physical Volumes Before and After VTS APM Release

The following example assumes that five physical backend drives are available for premigration during the writing of 25 logical volumes. The number of logical volumes grouped on the same physical volume will vary depending on factors including volume size and workload.

Representative Logical Volume Distribution Prior to VTS APM Release

PV1	LV1	LV6	LV11	LV16	LV21
PV2	LV2	LV7	LV12	LV17	LV22
PV3	LV3	LV8	LV13	LV18	LV19
PV4	LV4	LV9	LV14	LV19	LV20
PV5	LV5	LV10	LV15	LV20	LV21

Representative Logical Volume Distribution After VTS APM Release

PV1	LV1	LV2	LV3	LV4	LV5
PV2	LV6	LV7	LV8	LV9	LV10
PV3	LV11	LV12	LV13	LV14	LV15
PV4	LV16	LV17	LV18	LV19	LV20
PV5	LV21	LV22	LV23	LV24	LV25

VTS FICON Performance at Distance

- Distance measurements were made on a model B20 using a pair of IBM 2029 Fibre Savers, driven from a Freeway host processor.
- FICON supports distances of up to 100 kilometers, vs. 50 kilometers maximum for ESCON.
- Single-job performance
  - Less relative data rate loss (droop) than for ESCON at the same distance.
  - The larger the block size the less the droop, with 256K blocks showing almost no droop for FICON at 100 KM.
  - Single-job distance droop will be the same for Base and PAF FICON.
- Multi-job performance
  - 32K blocks and greater show negligible droop for up to 100 KM.

Chart 22: FICON
Single-job Write Performance vs. Distance vs. Blocksize

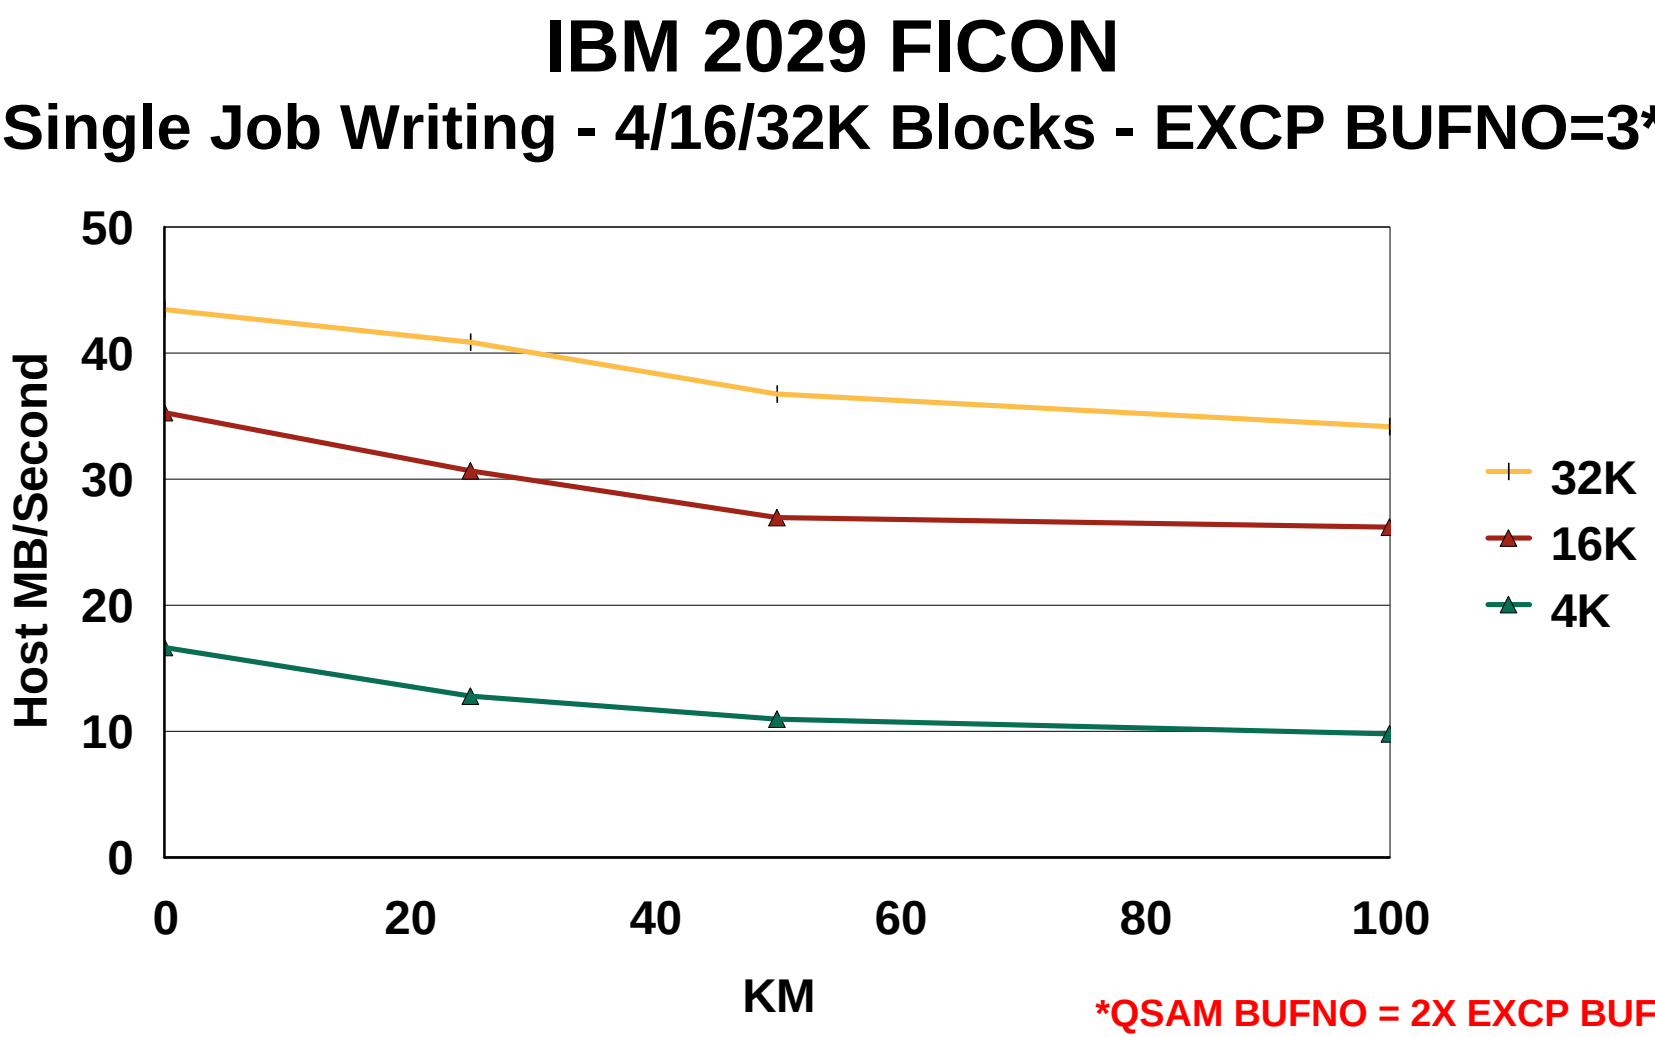

*QSAM BUFNO = 2X EXCP BUFNO

Chart 23: FICON
Single-job Read Performance vs. Distance Vs. Blocksize

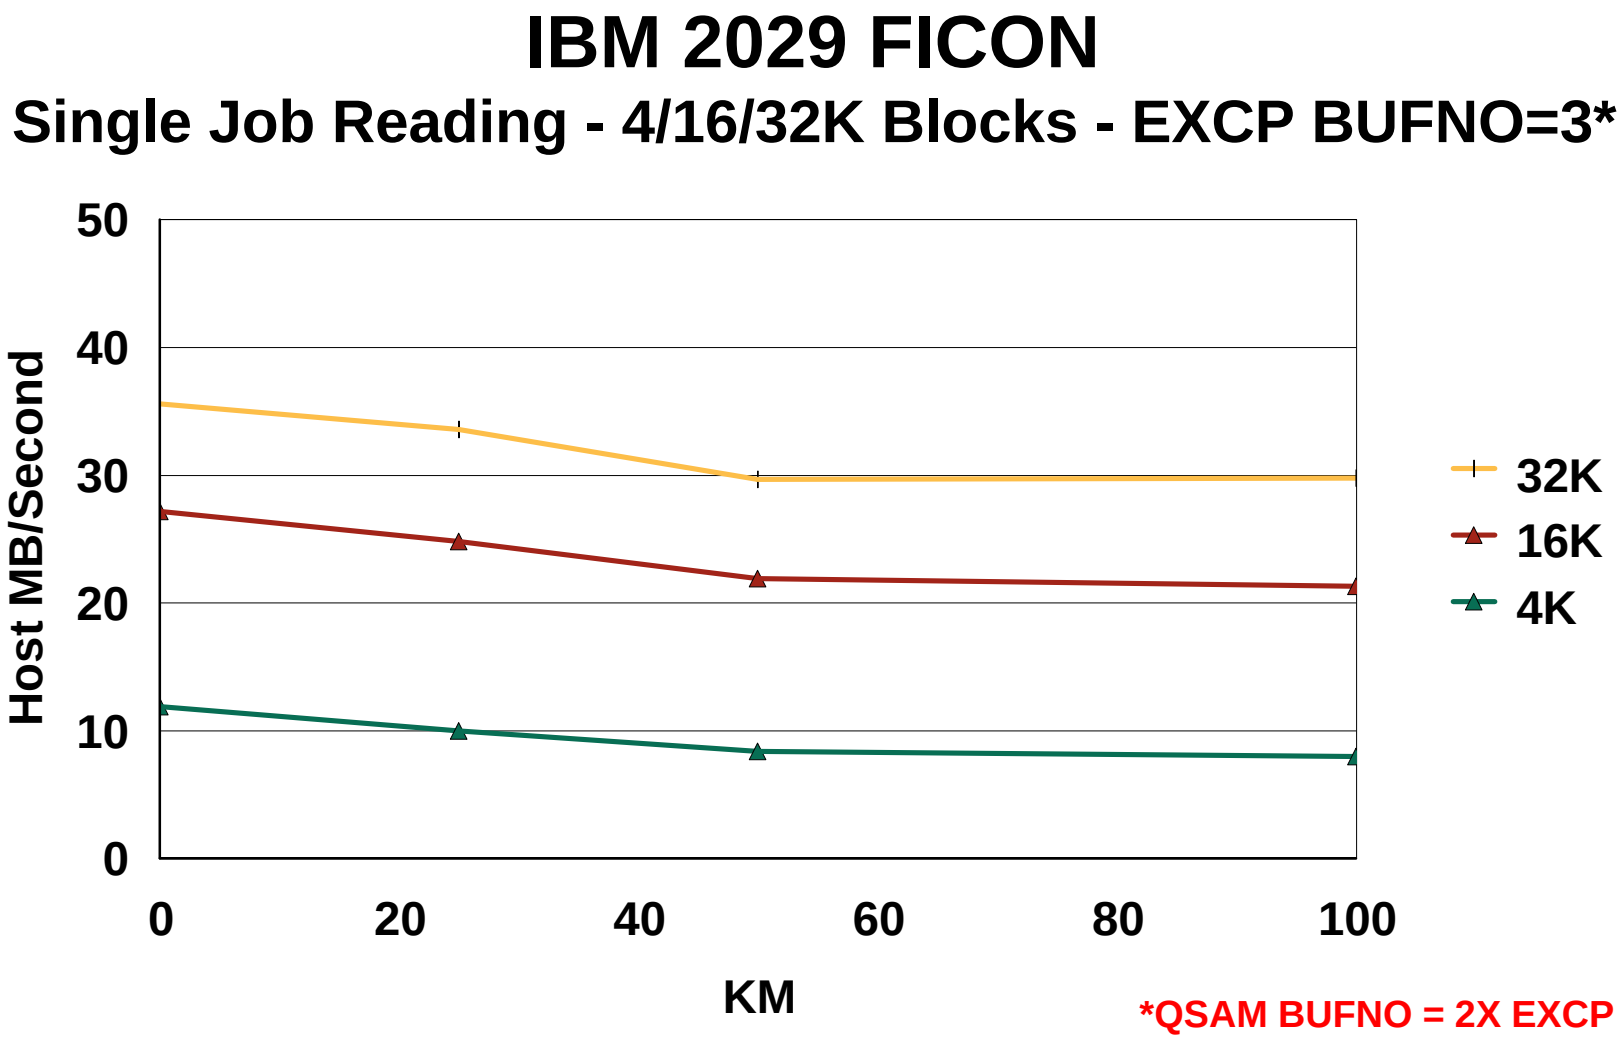

*QSAM BUFNO = 2X EXCP BUFNO

Chart 24: ESCON / FICON
Single-job Write Performance vs. Distance

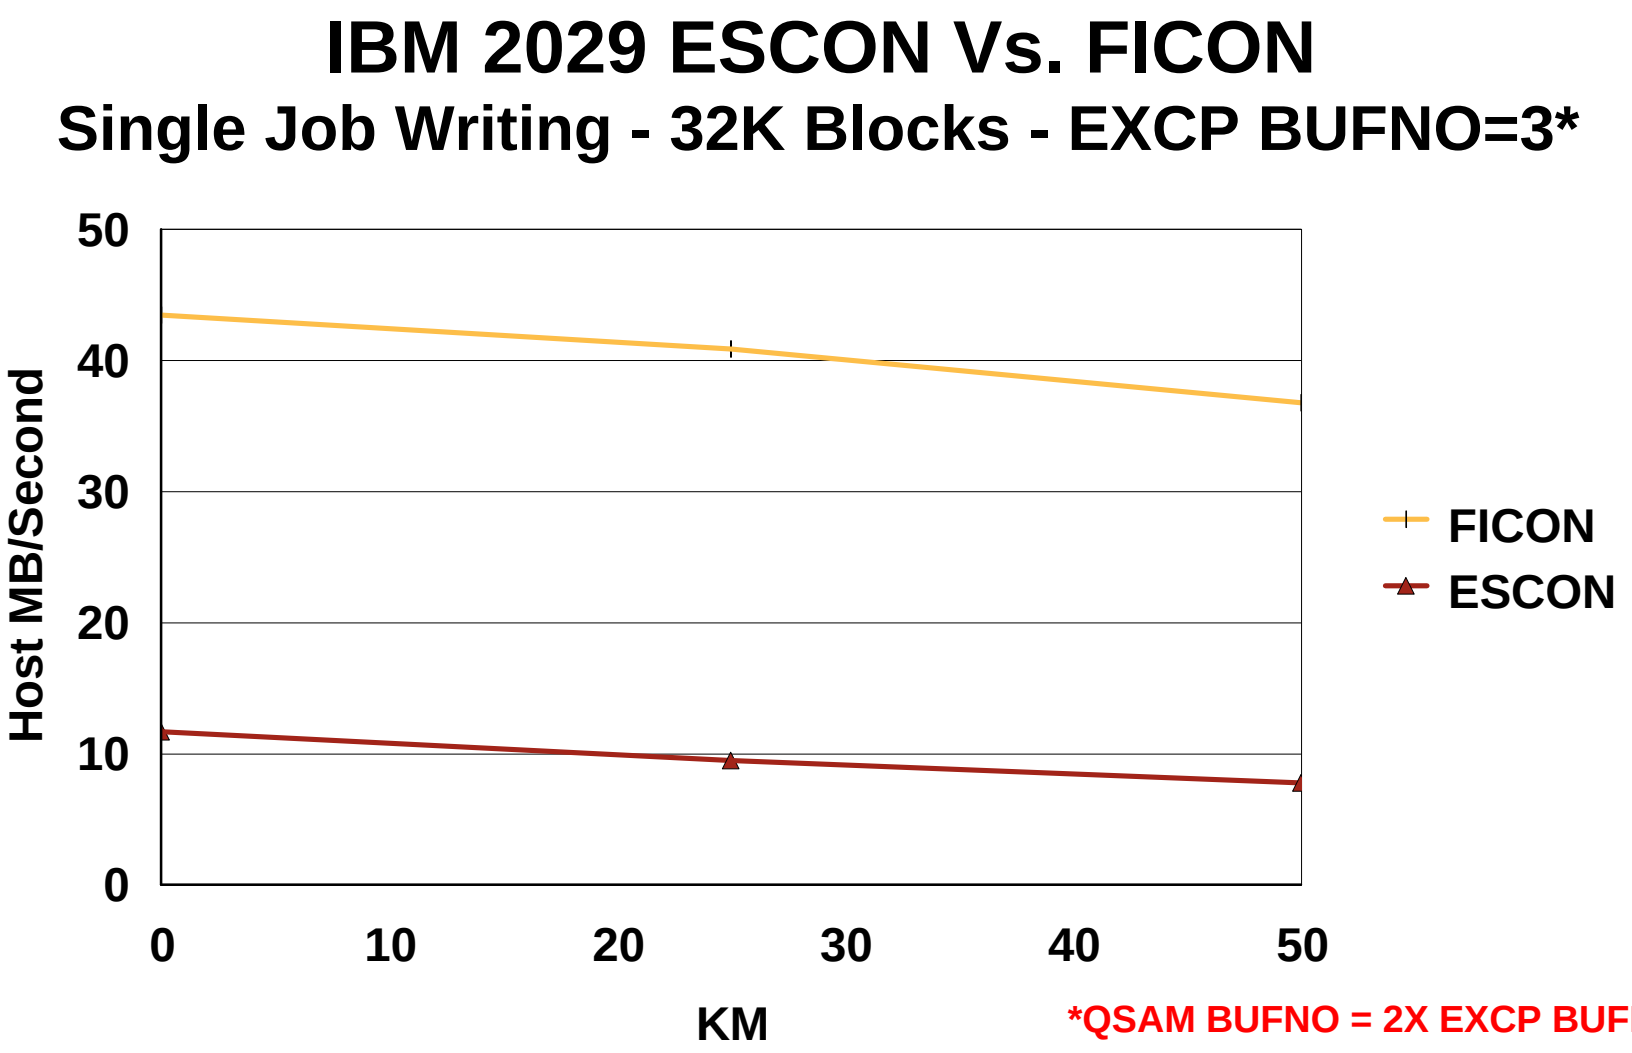

Chart 25: ESCON / FICON
Single-job Read Performance vs. Distance

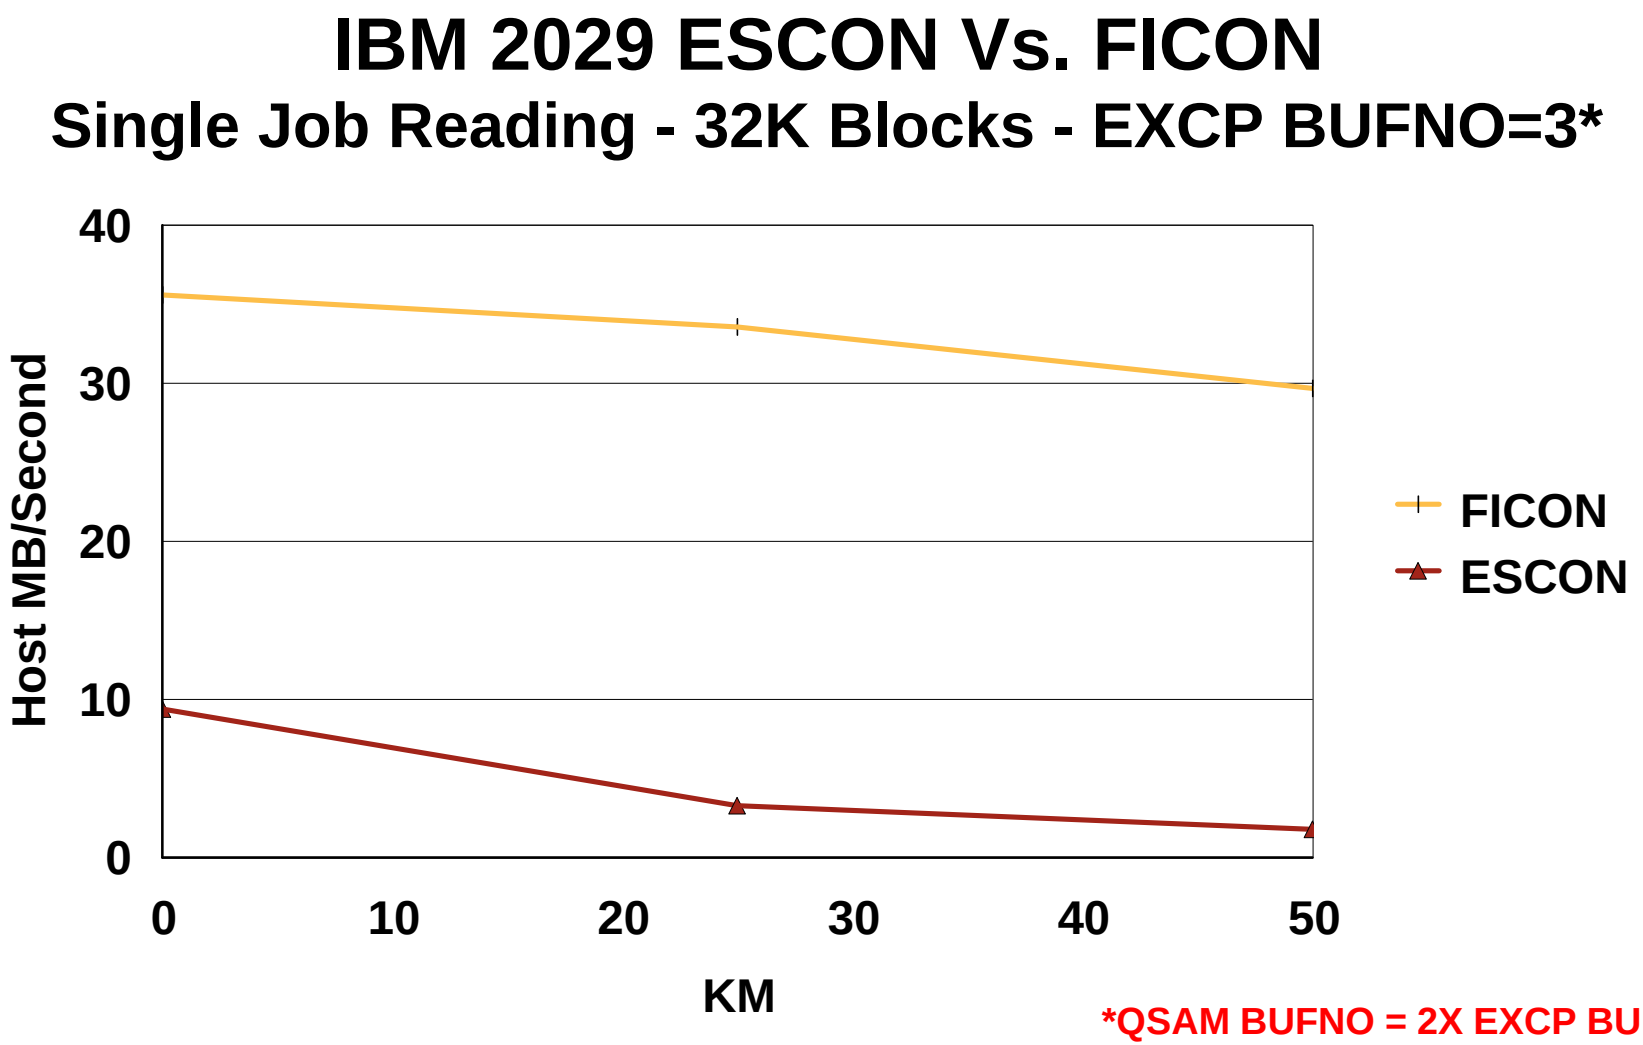

*QSAM BUFNO = 2X EXCP BUFNO

Advanced Policy Management

Performance Overview

Advanced Policy Management Overview

- Enables customer control of data management and placement in a VTS
  - Similar to existing DASD data management and placement controls
  - Extension to DFSMS storage constructs and automatic class section routines
- Combination of software and hardware functions
  - SPE to existing DFSMS/MVS system support levels
  - New levels of VTS and Library Manager LIC
- Activated via Feature Code
  - Feature Code 4001, Advanced Policy Management
    - Upgrade from Feature Code 4000, Advanced Functions
  - Enabled by customer through the Library Manager's controls

Advanced Policy Management Functions

- Physical Volume Pooling
  - Associates logical volumes with a set of physical volumes
  - Controlled through the Storage Group construct
- Selective Dual Copy
  - Creates two copies of a logical volume on different physical volume pools
  - Controlled through the Management Class construct
- Tape Volume Cache Management
  - Control of which logical volumes have preference to be kept in cache
  - Controlled through the Storage Class construct
- PtP Copy Mode Control
  - Control of how the copy of a logical volume between VTSs is performed
  - Controlled through the Management Class construct

APM Functions Performance Measurements

- Neither volume pooling nor selective dual copy result in a decrease in peak data rate.
- For pooling:
  - With just one pool there is no reduction in sustained data rate.
  - With a number of pools less than the number of physical drives (per Bxx for a PtP), there may be up to a 10% reduction in sustained data rate.
  - With a number of pools equal to or greater than the number of physical drives, there may be up to a 15% reduction in sustained data rate.
- For dual copy:
  - With 100% dual copy there may be up to a 50% reduction in sustained data rate.

Performance Benefits of Volume Pooling for Export Operations

- Physical volume pooling can reduce logical volume export time by directing logical volumes to be exported to the same physical volume(s).
- The following chart shows measured export times for various combinations of logical and physical volumes:
  - Export 10 logical volumes from 10 physical volumes
  - Export 10 logical volumes from one physical volume
  - Export 50 logical volumes from 50 physical volumes
  - Export 50 logical volumes from 10 physical volumes
  - Export 50 logical volumes from one physical volume
- These measurements were made from a host G6 processor on a model B20 VTS with 12 backend drives. Each volume was 800MB prior to 2.66:1 data compression. Physical cartridges were "J" cartridges.

Chart 26: B20
Export Times

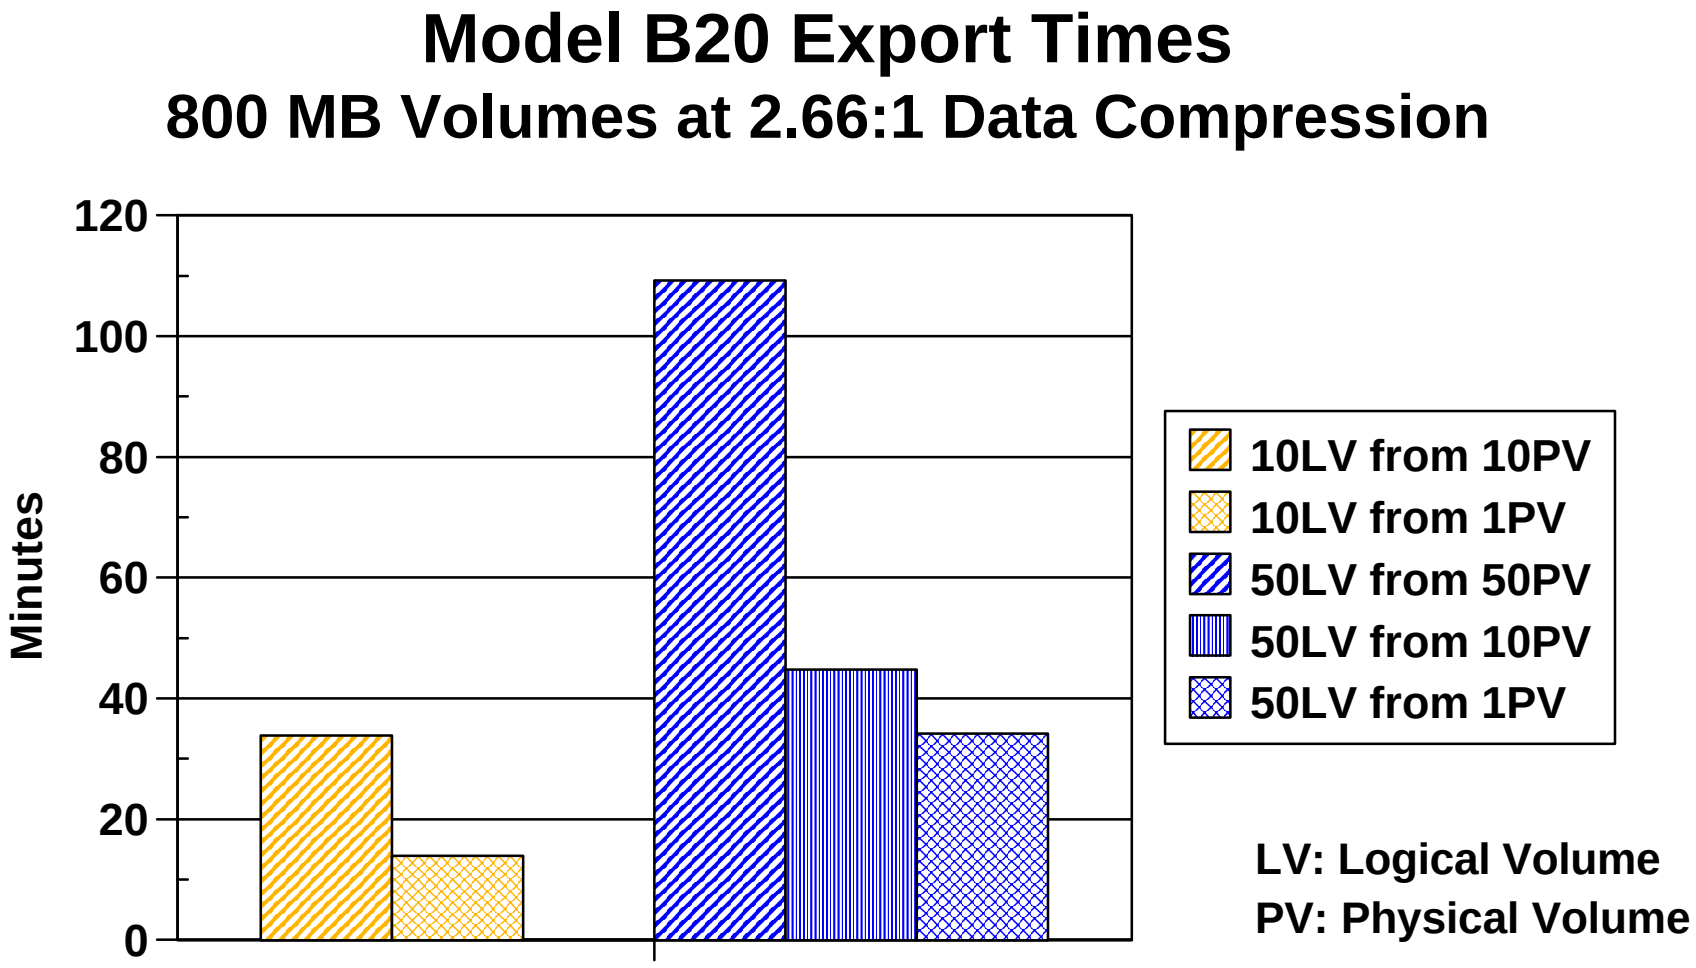

Logical Mount Performance Improvement

- This VTS release increases the maximum logical mount rate by 30%, to over 3,000 logical mounts per hour for a 3494 library with the newest Library Manager PC. The current version of this PC has been shipping with new 3494s since mid-December, 2000, and is available as an upgrade for older 3494s.

- Fast ready and read-hit mount times will accordingly also be reduced, with the amount of reduction depending on the rate and pattern of the logical mount requests.

End of Document