SNIA ESF | ETHERNET STORAGE

STORAGE PERFORMANCE BENCHMARKING: PART 4 – FILE COMPONENTS

Ken Cantrell, NetApp
Mark Rogov, Dell EMC

October 20, 2016

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

About SNIA

SNIA ESF | ETHERNET STORAGE

SNIA-At-A-Glance

160
unique member companies

3,500
active contributing members

50,000
IT end users & storage pros worldwide

Learn more: **snia.org/technical**

@SNIA

About The Speakers

Ken Cantrell
NetApp
Manager Perf Engineering
@kencantrelljr

Mark Rogov
Dell EMC
Advisory Systems Engineer
@rogovmark

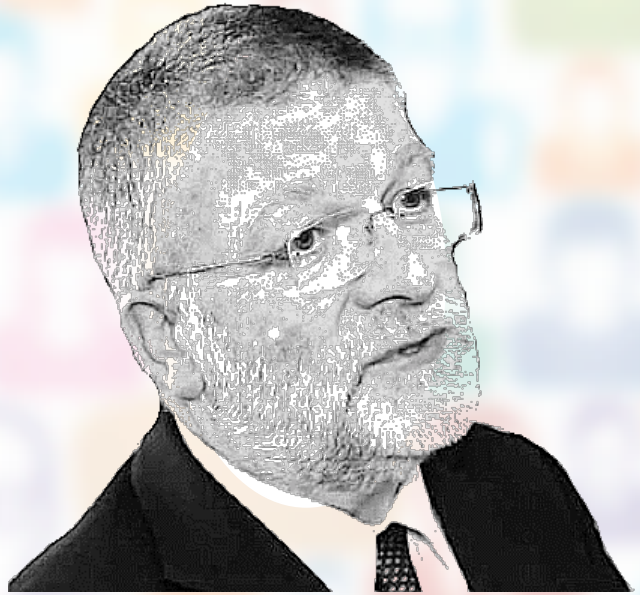

Alex McDonald
NetApp
Vice Chair of the SNIA ESF
@alextangent

Storage Performance Benchmarking

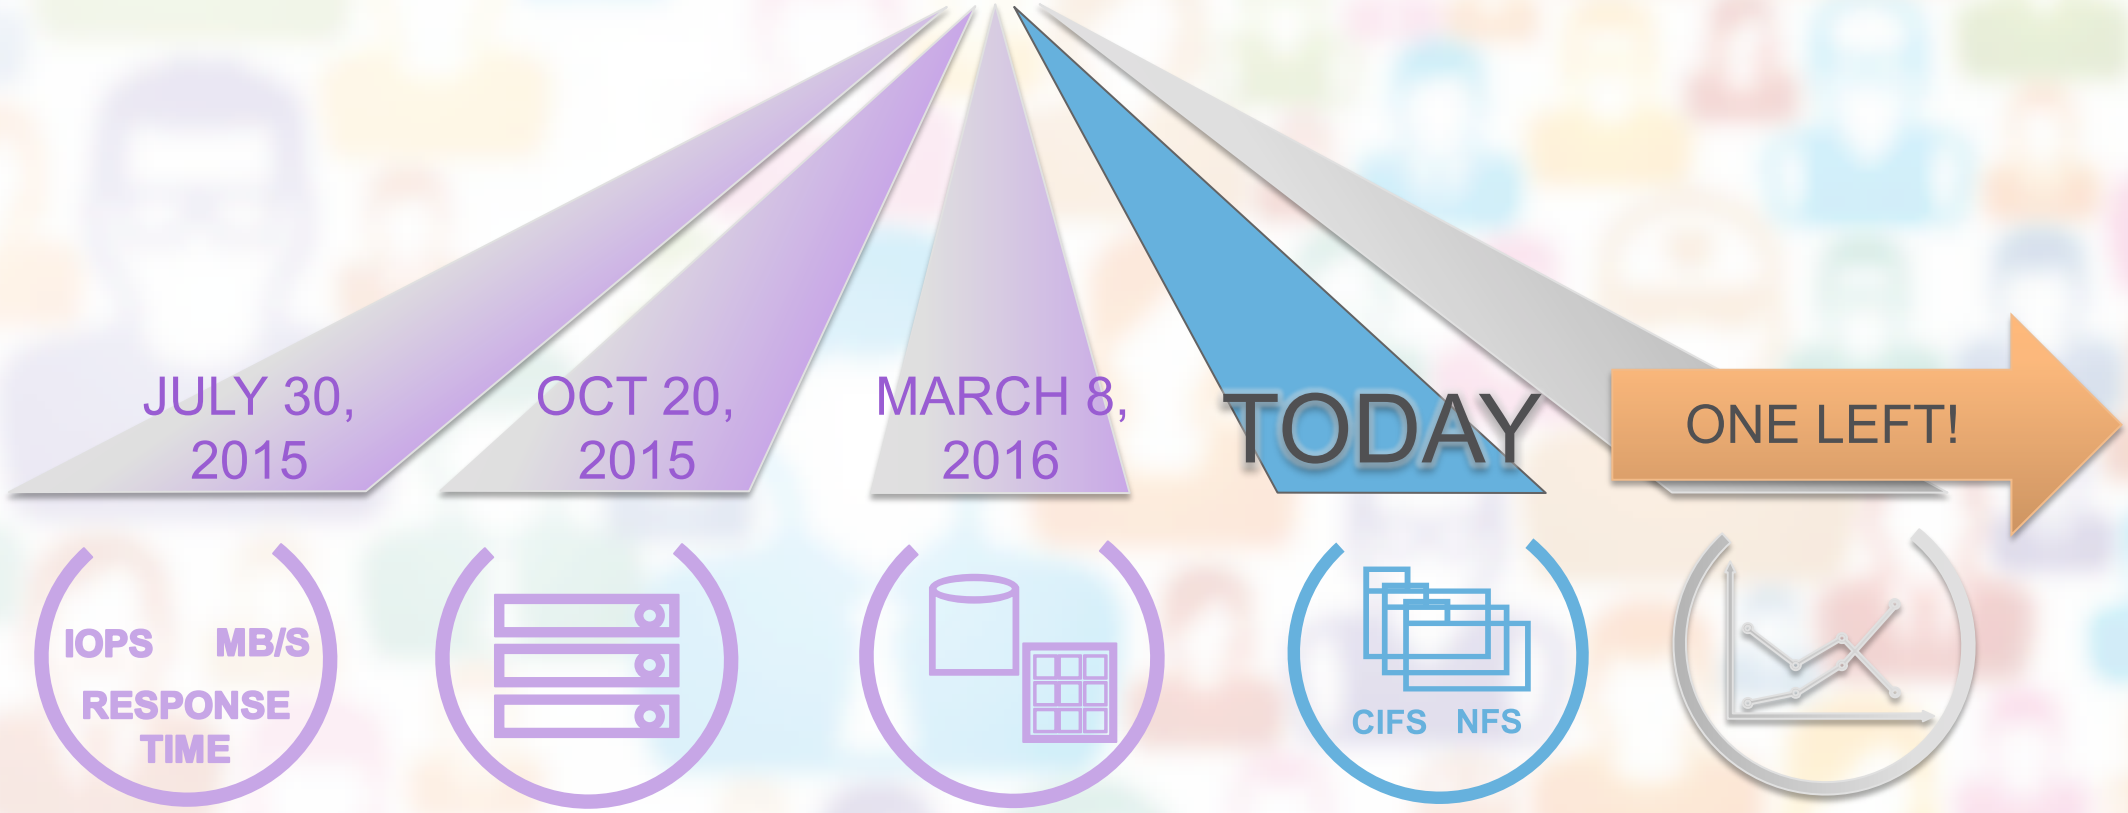

Session 1 – Terminology and Context

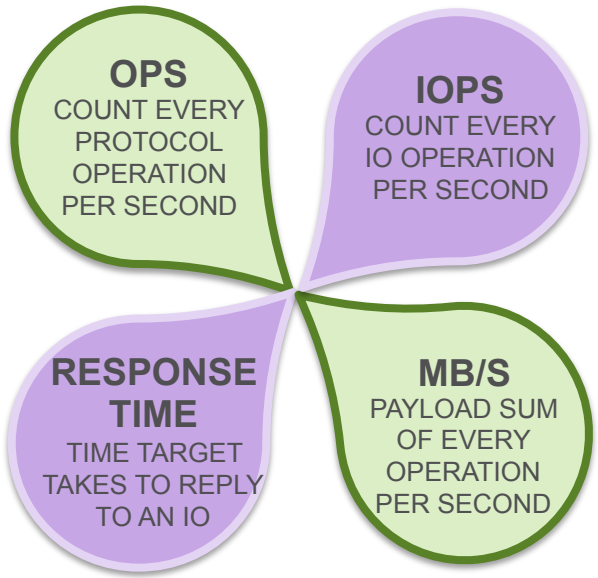

TERMINOLOGY AND BASIC CONCEPTS

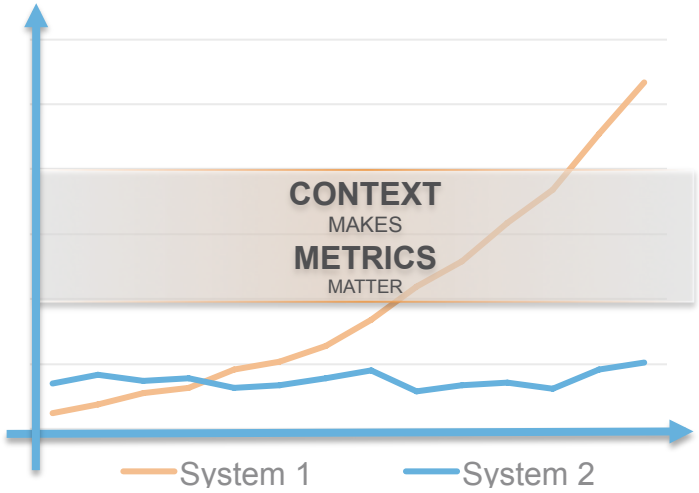

GRAPH FUN

Session 2 – The Slowest Component Matters Most

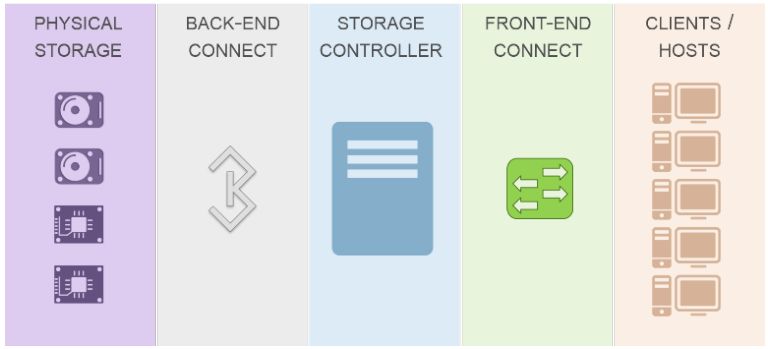

SLOW COMPONENT MATTERS MOST

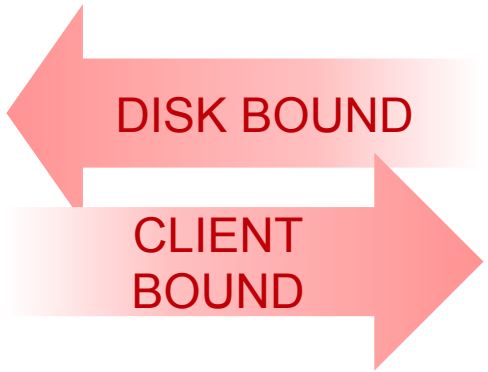

BOTTLENECKS ALWAYS EXIST

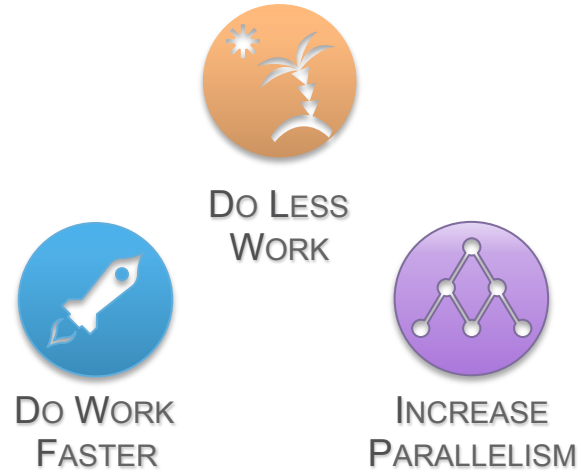

3 PERFORMANCE PRINCIPLES

Session 3 – Block Is The Foundation

RANDOM VS. SEQUENTIAL WORKLOADS

FLASH VS DISK DRIVE

RAID PERFORMANCE IMPLICATIONS

Today: File & Filesystem Performance

COUNTS

245

FILESYSTEMS

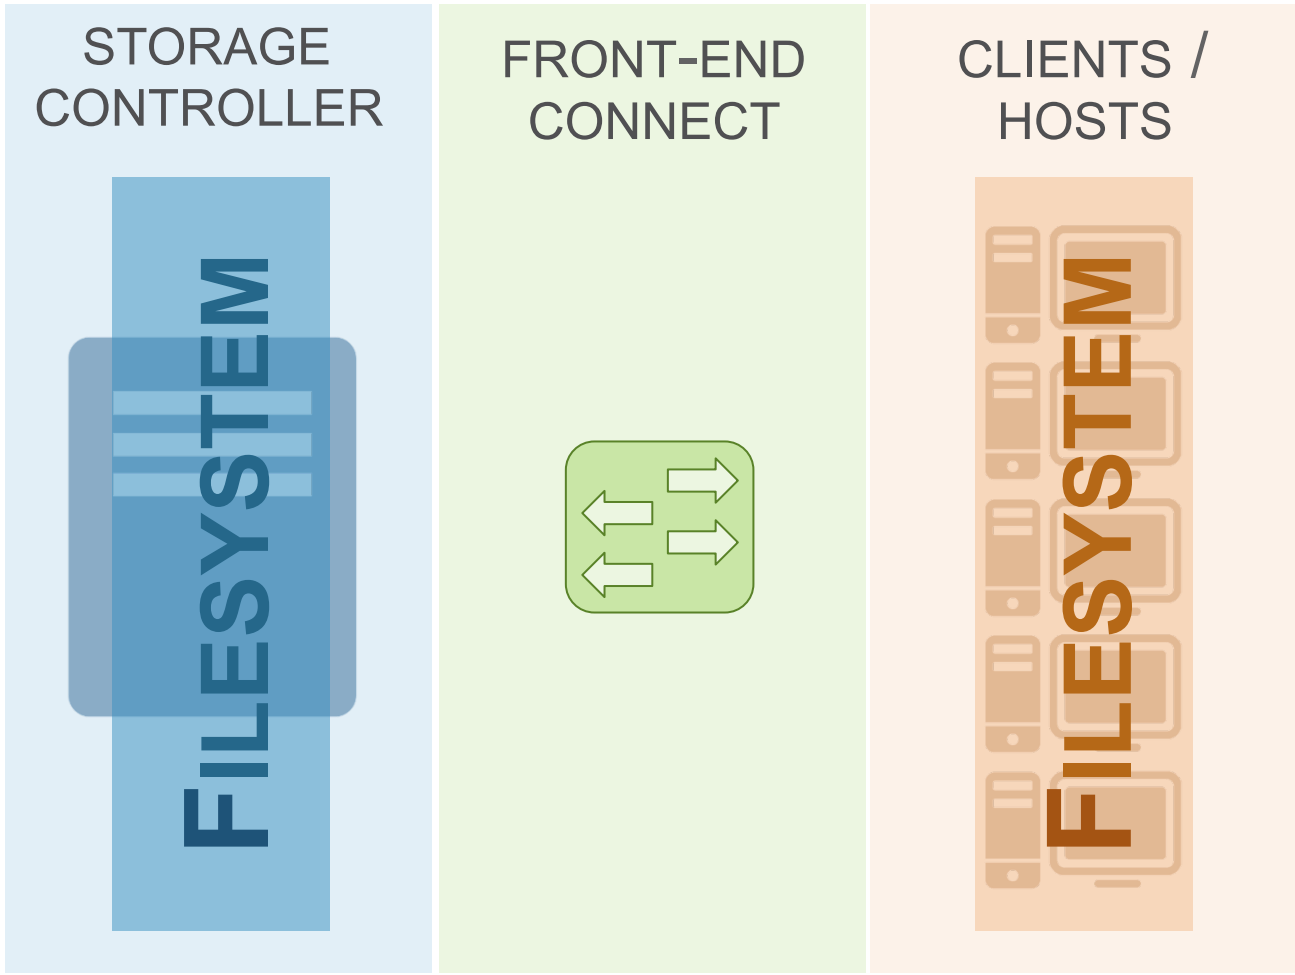

Agenda

INTRO	INTRODUCTION
FILESYSTEM	FILESYSTEM LAYOUT FUNDAMENTALS
FS OPS	WHAT ARE FILESYSTEM OPERATIONS?
FILE PERF	PERFORMANCE IN A FILE ORIENTED WORLD
FUN	GRAPH FUN WITH PERFORMANCE
END	SUMMARY

What Is A Filesystem?

Term “Filesystem” was coined in 1964

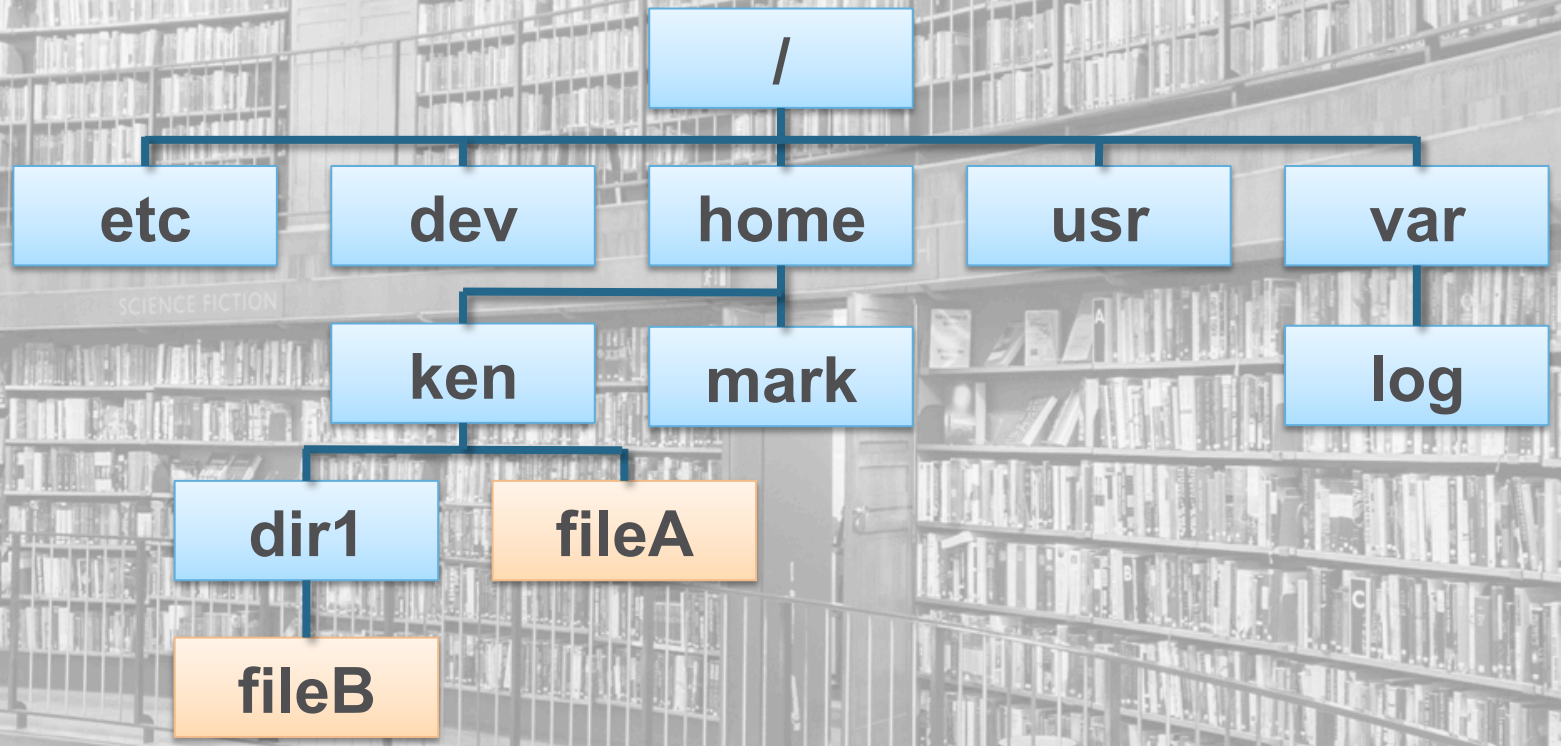

Filesystem Metadata

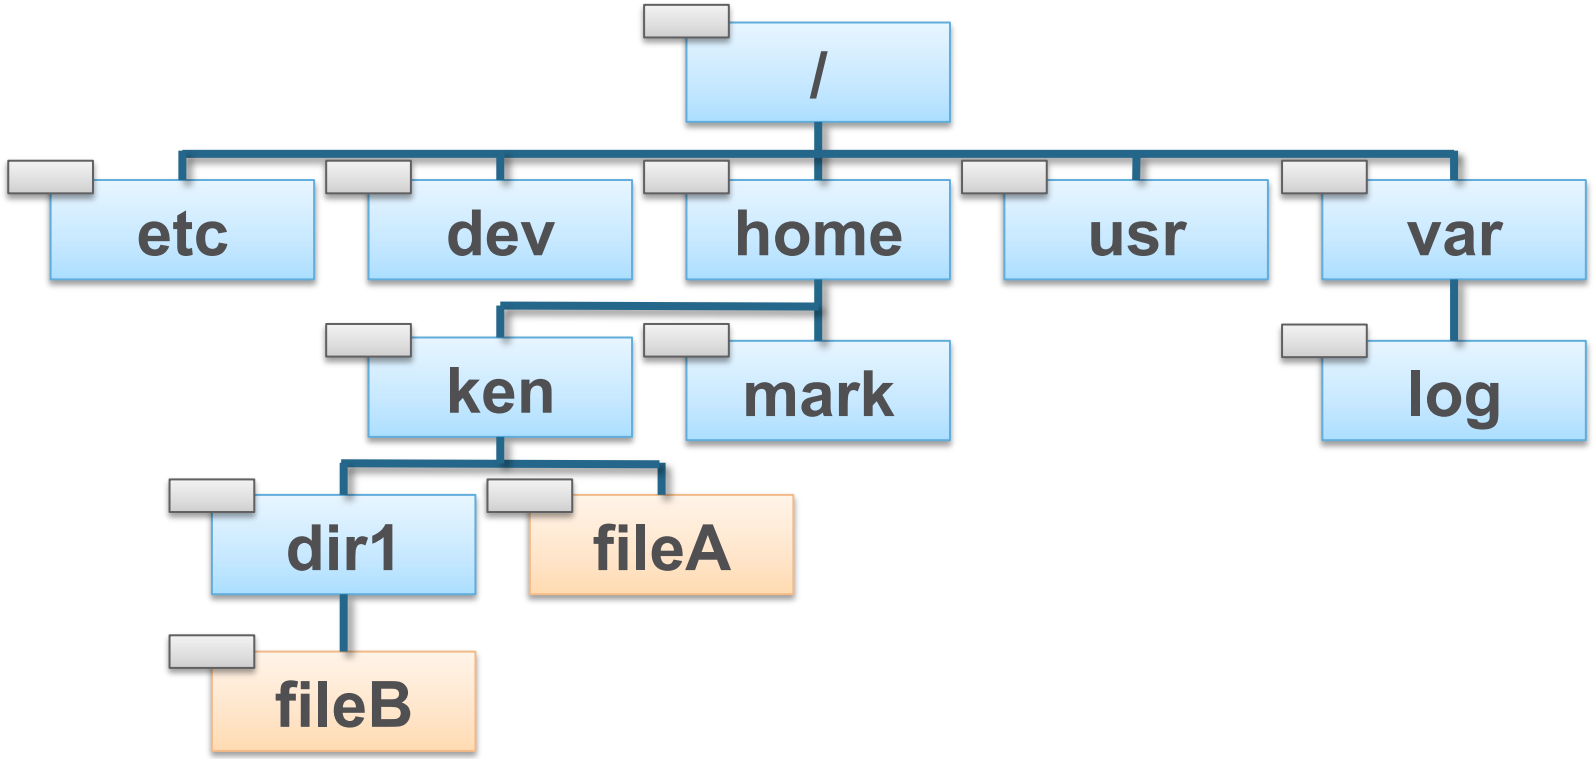

METADATA

- NAME
- PERMISSIONS
- ACCESS TIMES

… AND MORE …

- DATA BLOCK POINTERS

Filesystem Basic Metadata Structure

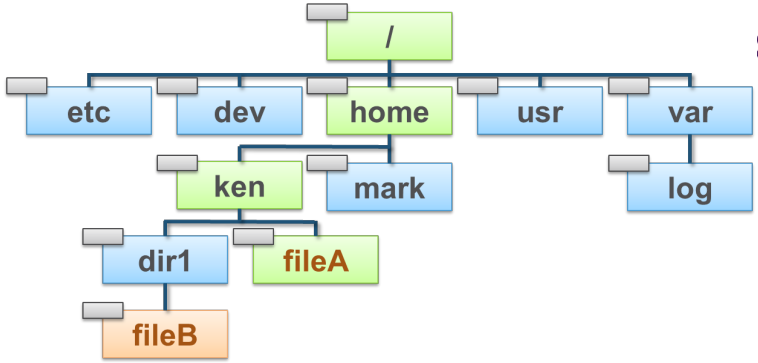

SNIA ESF | ETHERNET STORAGE

/home/ken/fileA

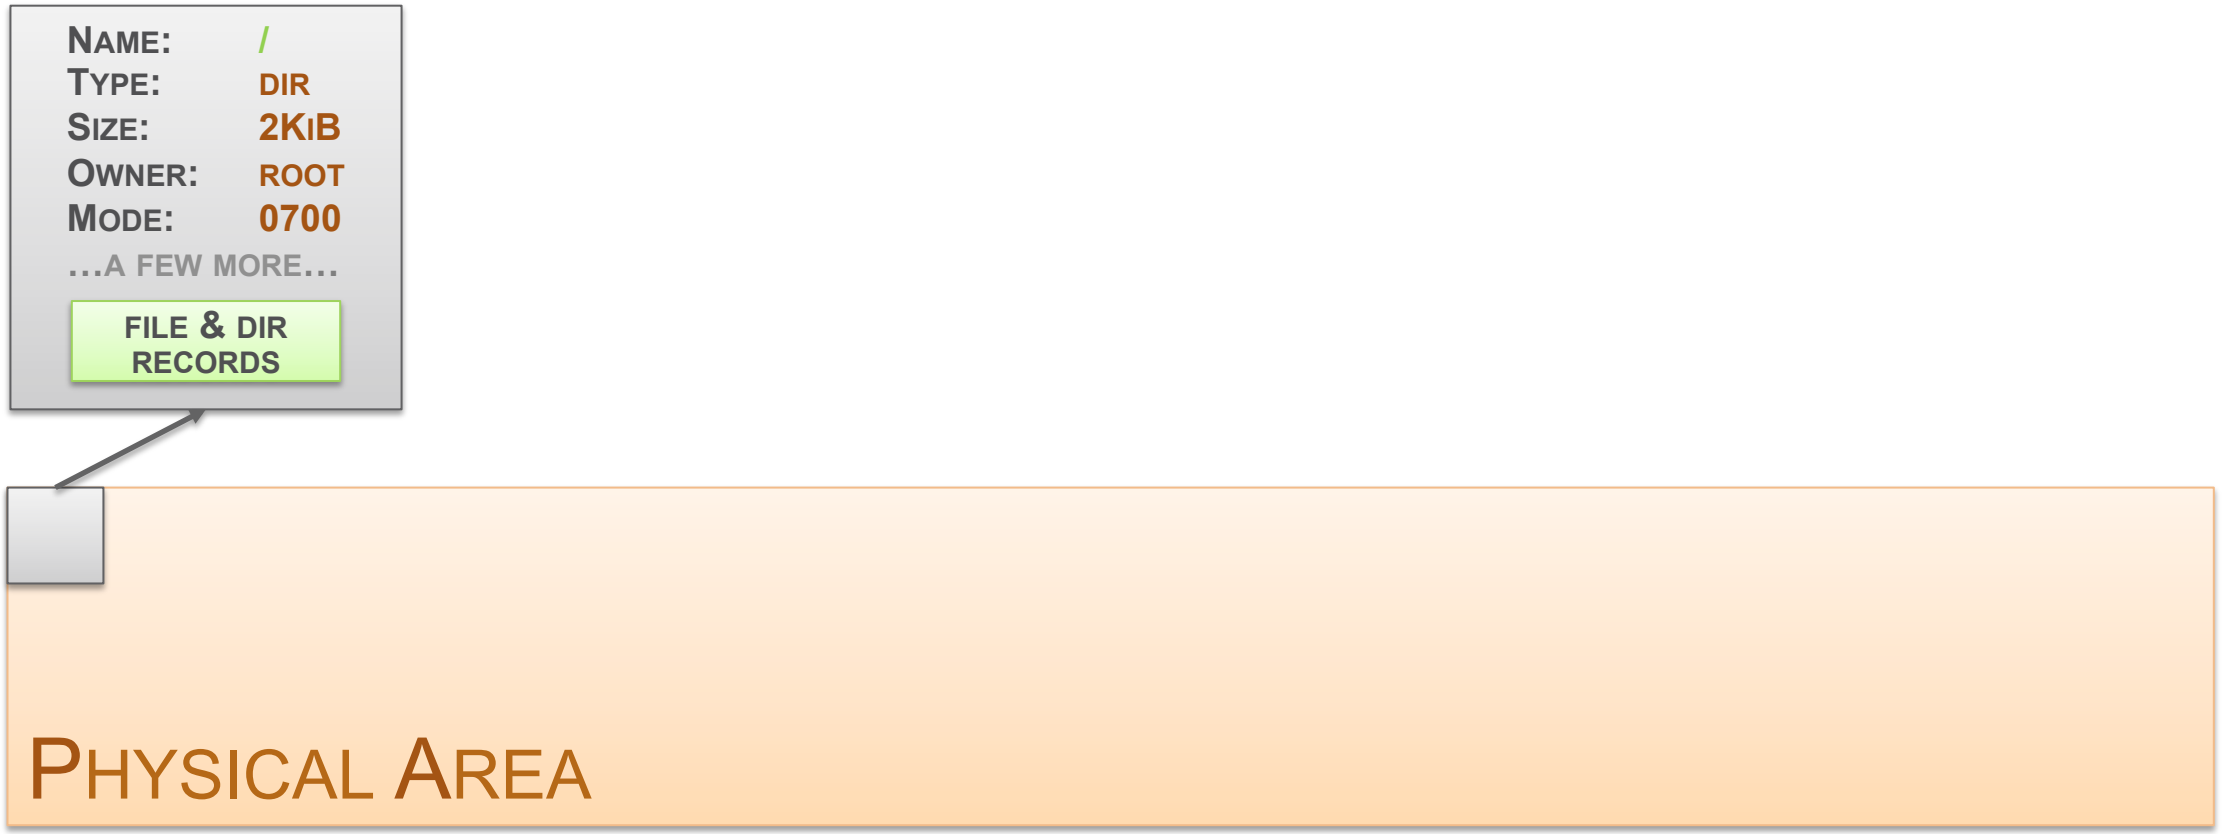

Filesystem Basic Metadata Structure

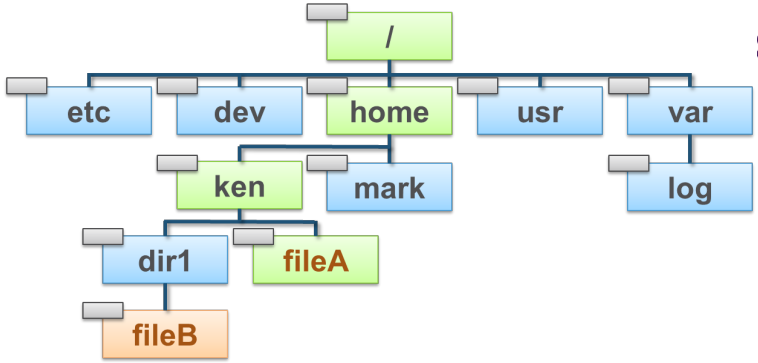

SNIA ESF | ETHERNET STORAGE

/home/ken/fileA

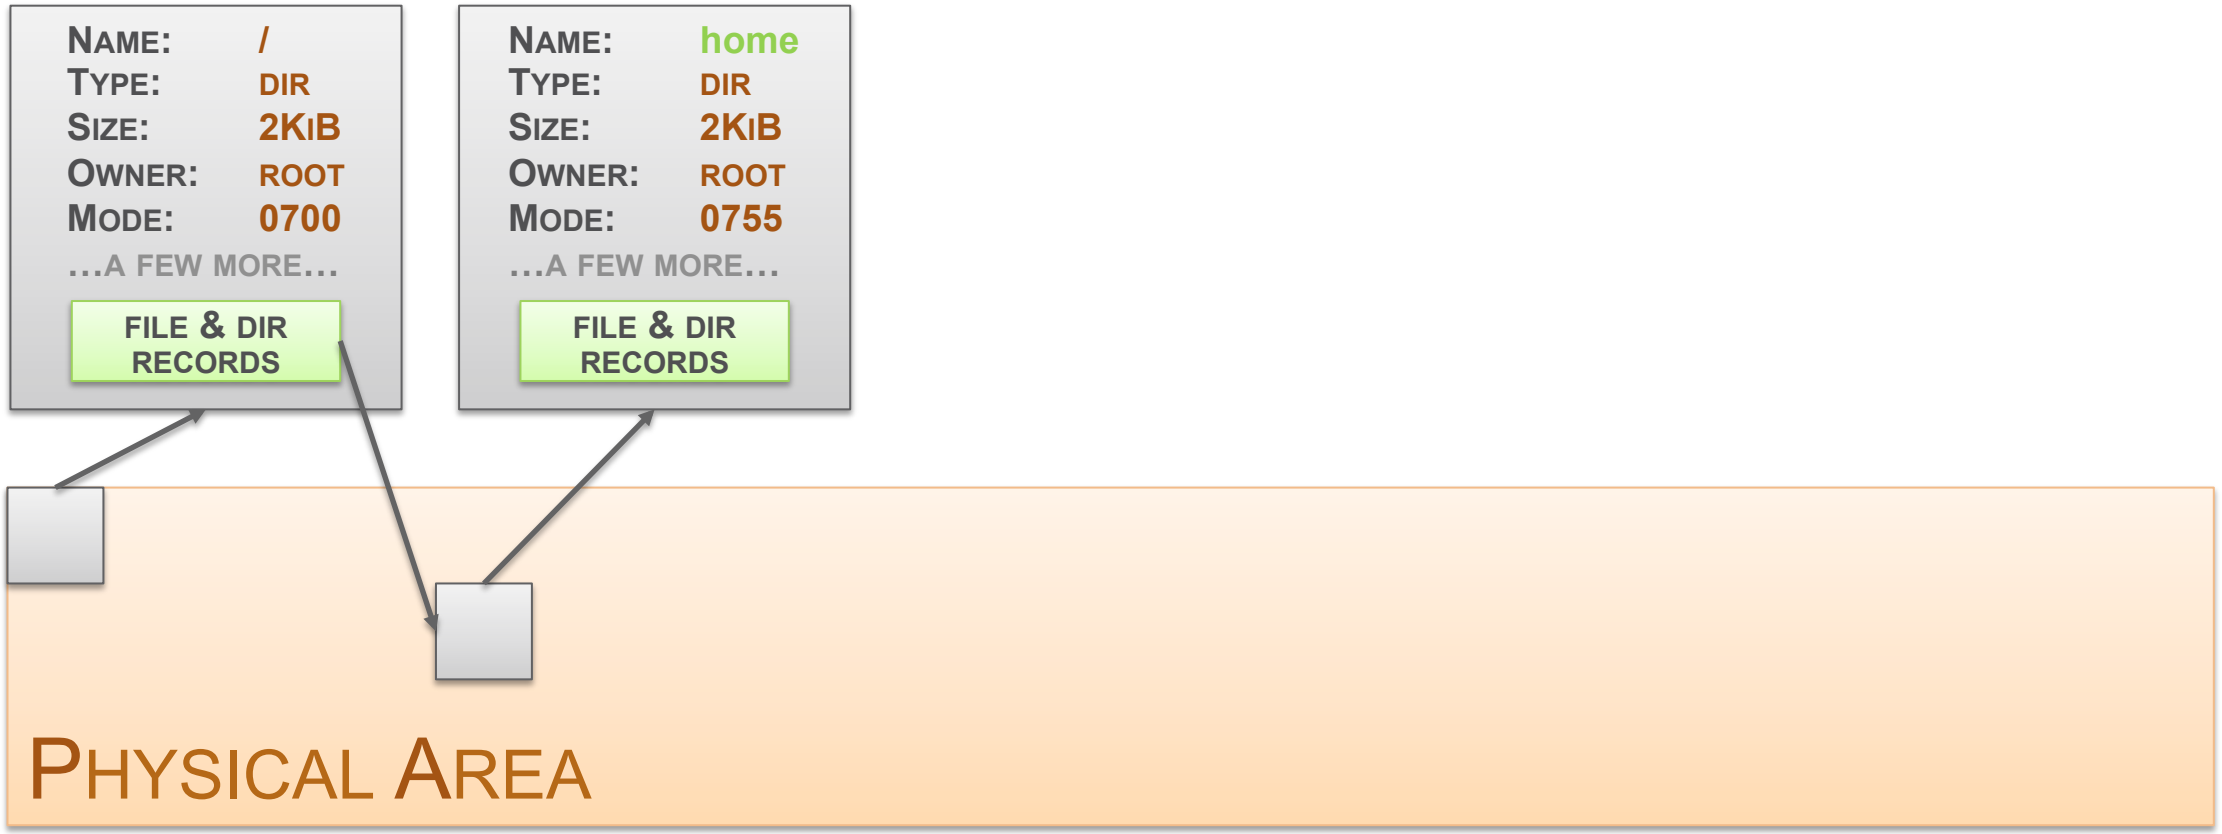

Filesystem Basic Metadata Structure

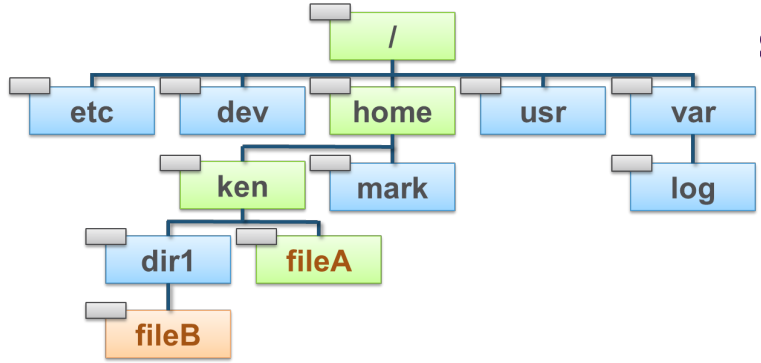

/home/ken/fileA

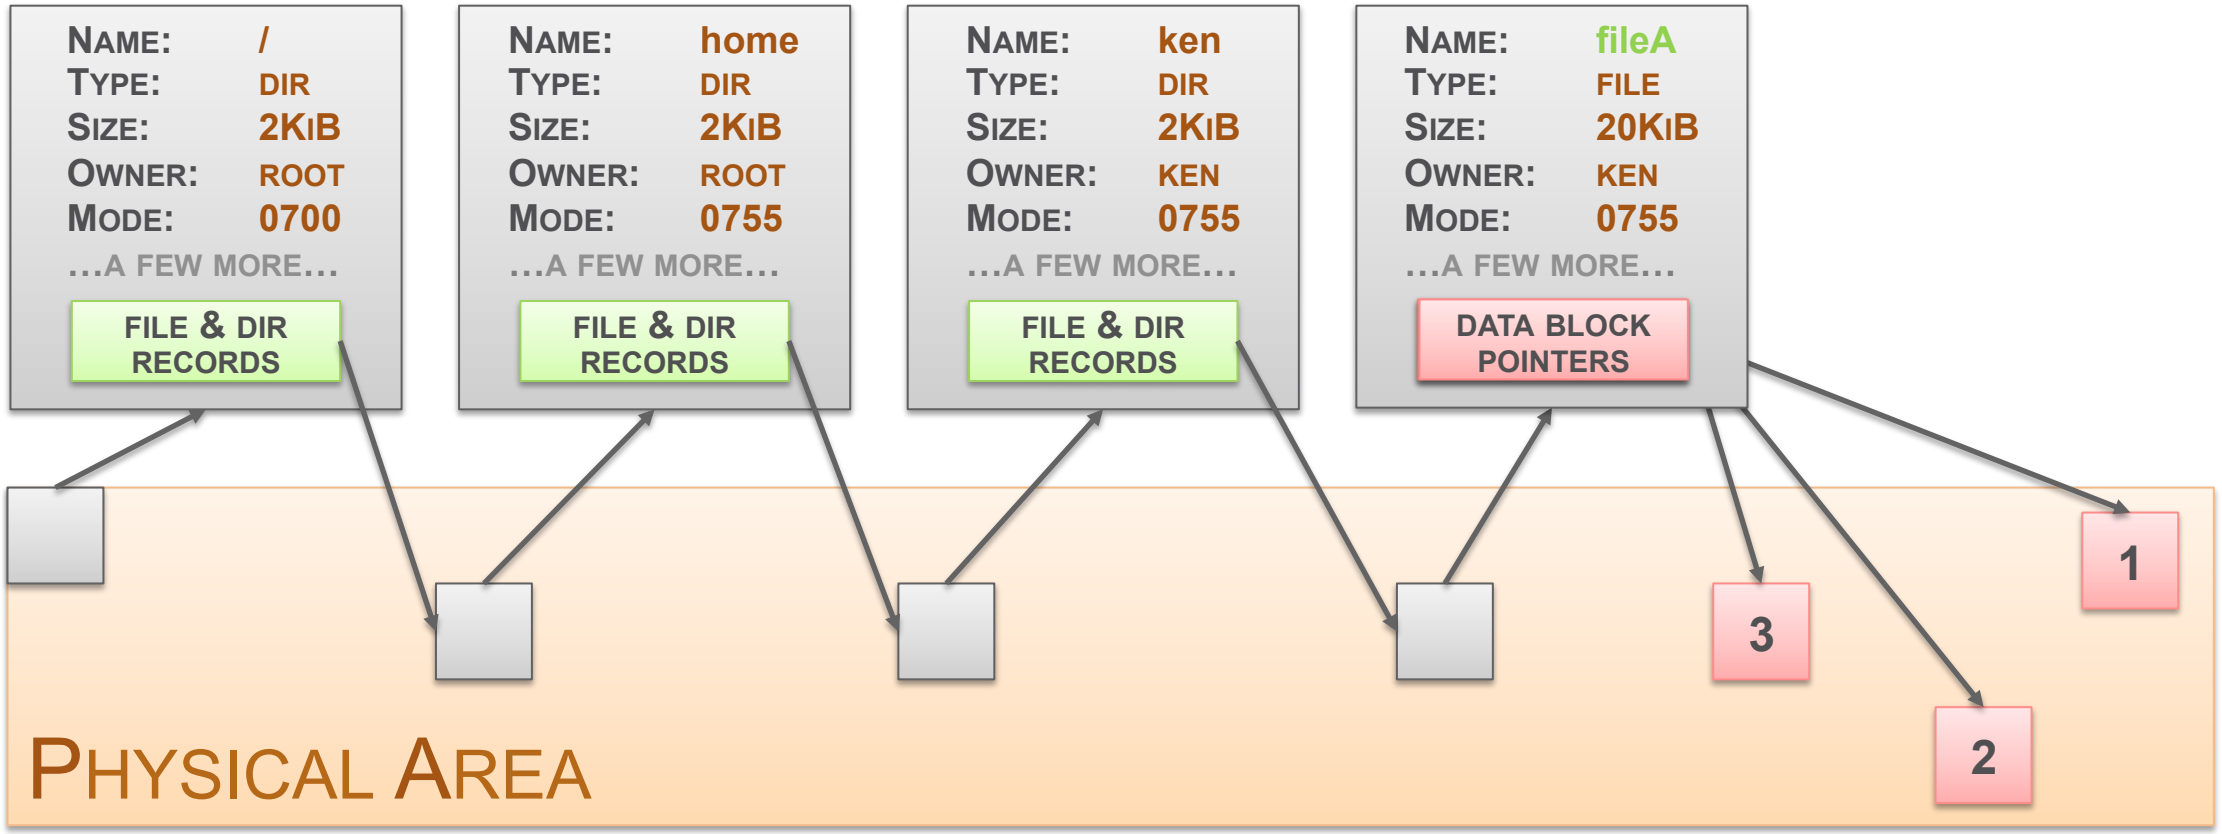

Filesystem Metadata

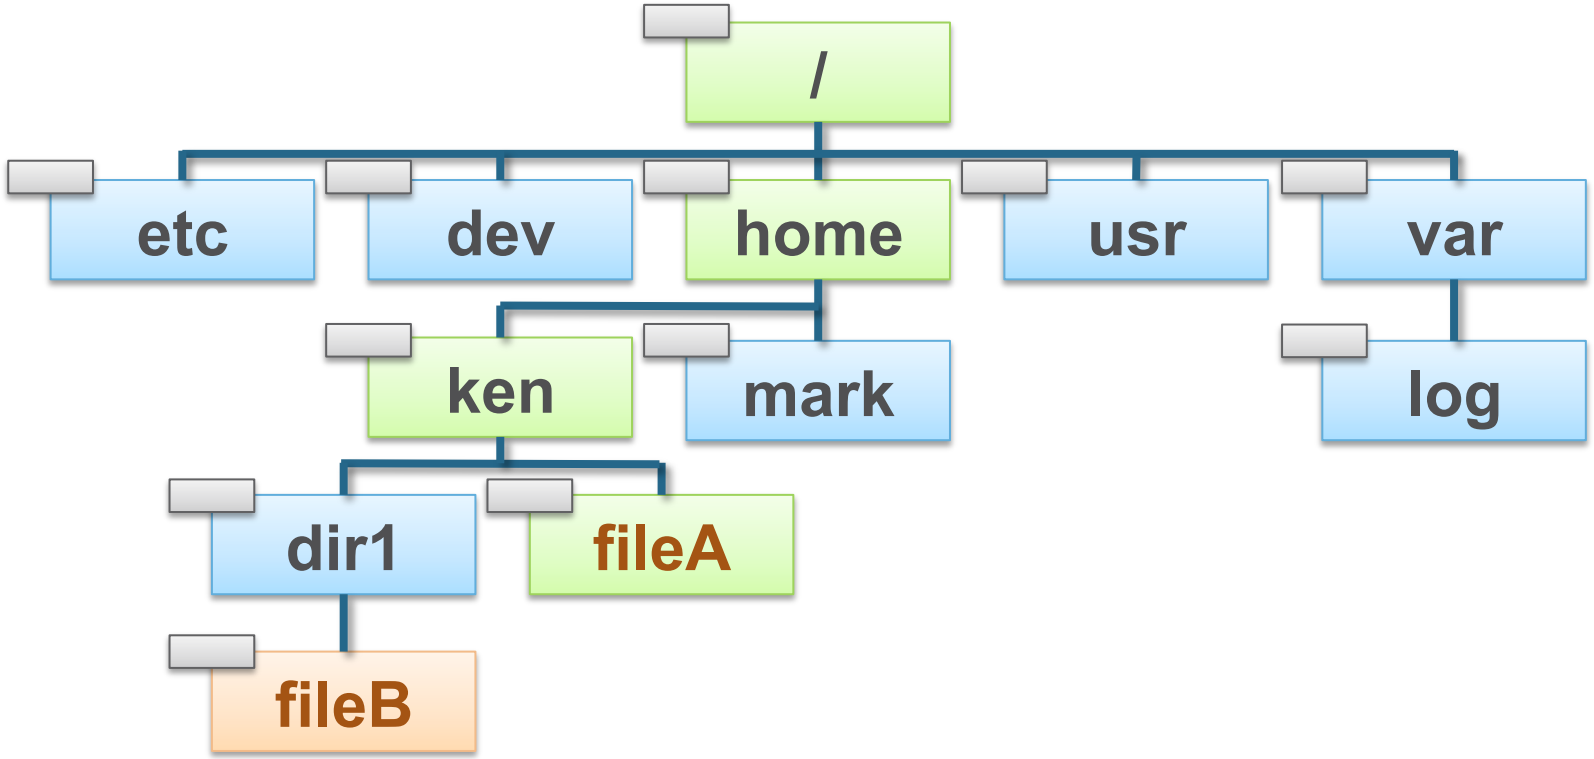

METADATA

- NAME
- PERMISSIONS
- ACCESS TIMES

… AND MORE …

- DATA BLOCK POINTERS

Filesystem Metadata

FILESYSTEMS PROVIDE STRUCTURE

Agenda

INTRO	INTRODUCTION
FILESYSTEM	FILESYSTEM LAYOUT FUNDAMENTALS
FS OPS	**WHAT ARE FILESYSTEM OPERATIONS?**
FILE PERF	PERFORMANCE IN A FILE ORIENTED WORLD
FUN	GRAPH FUN WITH PERFORMANCE
END	SUMMARY

Local File Operations

User Action Converted Into Filesystem Operations

```
$ echo "test" > data.txt
```

STAT OPEN WRITE CLOSE

Local File Operations

SNIA ESF | ETHERNET STORAGE

User Action Converted Into Filesystem Operations

```
$ echo "test" > data.txt
```

STAT OPEN WRITE CLOSE

```
$ cat data.txt
```

STAT OPEN READ CLOSE

Local File Operations

SNIA ESF | ETHERNET STORAGE

User Action Converted Into Filesystem Operations

```
$ echo "test" > data.txt
```

STAT OPEN WRITE CLOSE

```
$ cat data.txt
```

STAT OPEN READ CLOSE

```
$ rm data.txt
```

STAT UNLINK

Local File Operations

Are There Only Six File Operations?

```
$ echo "test" > data.txt
```

STAT OPEN WRITE CLOSE

```
$ cat data.txt
```

STAT OPEN READ CLOSE

```
$ rm data.txt
```

STAT UNLINK

READ
WRITE
OPEN
STAT
UNLINK
CLOSE
…MORE…

COUNTS
108
FILESYSTEM CALLS

Local File Operations

Relationship Between Filesystem Operations And Block IO Operations

```
$ echo "test" > data.txt
```

STAT OPEN WRITE CLOSE

```
$ cat data.txt
```

STAT OPEN READ CLOSE

```
$ rm data.txt
```

STAT UNLINK

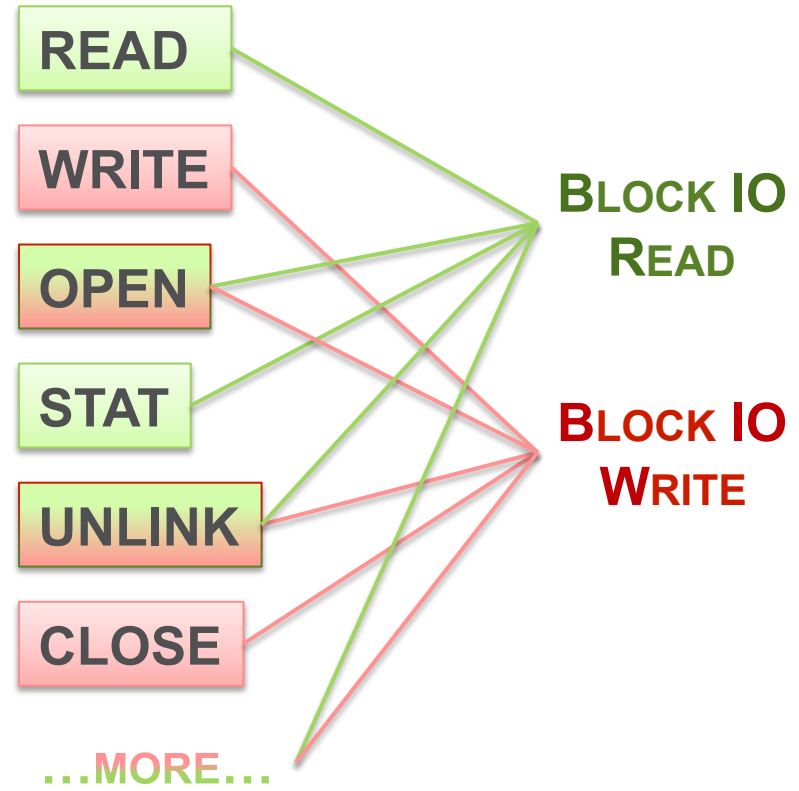

NAS File Operations

Create A New File Containing Just Three Spaces

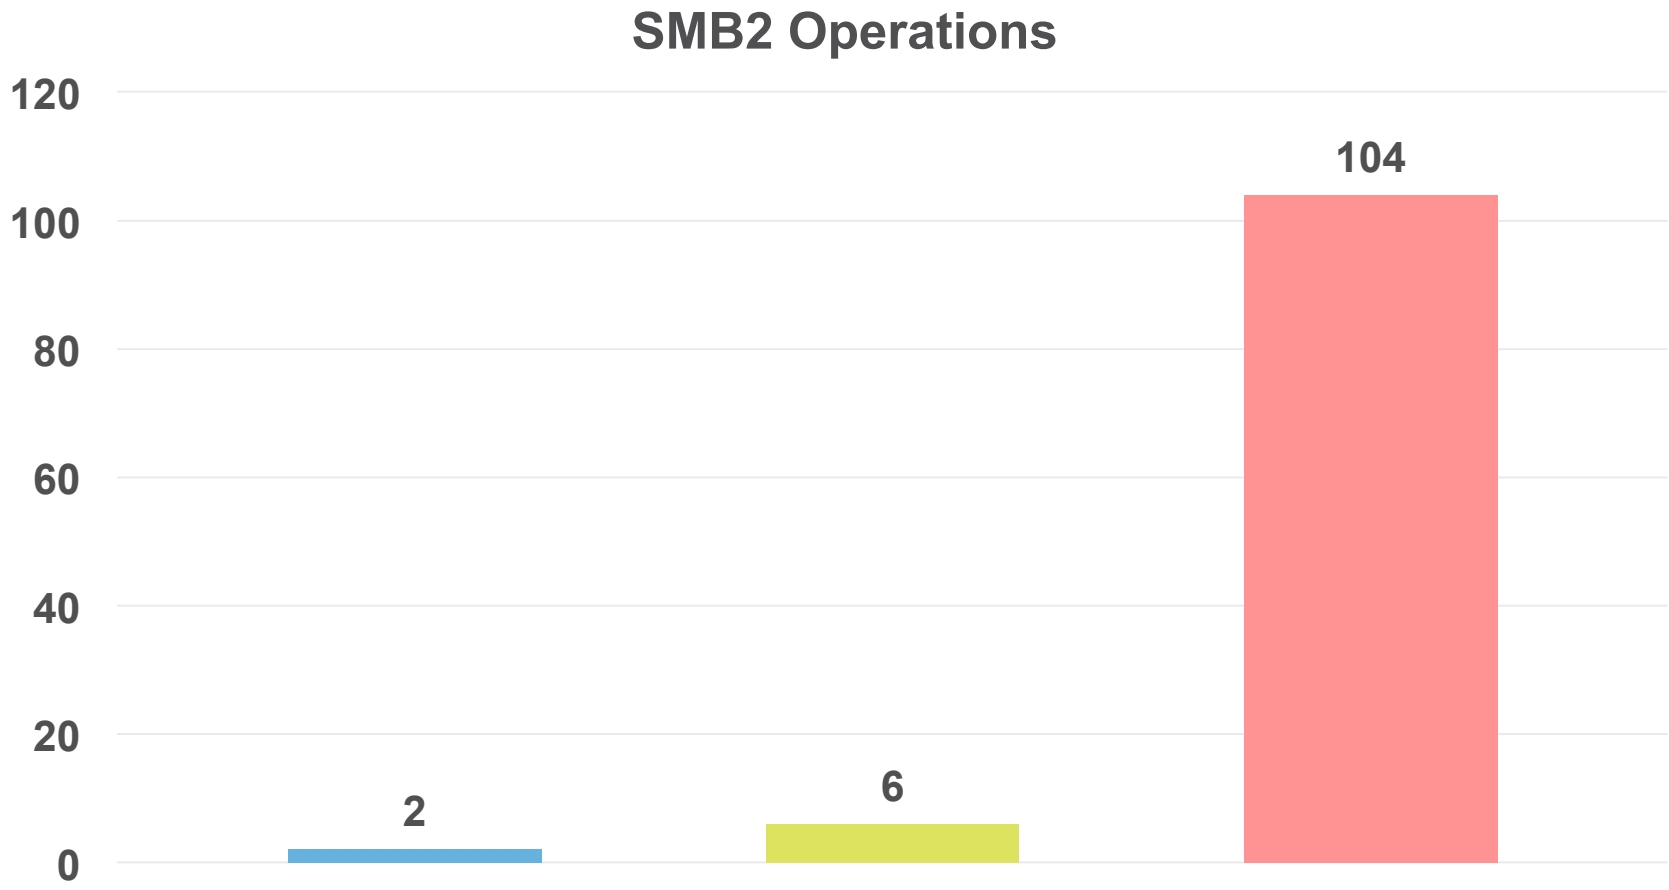

NAS File Operations

Create A New File Containing Just Three Spaces

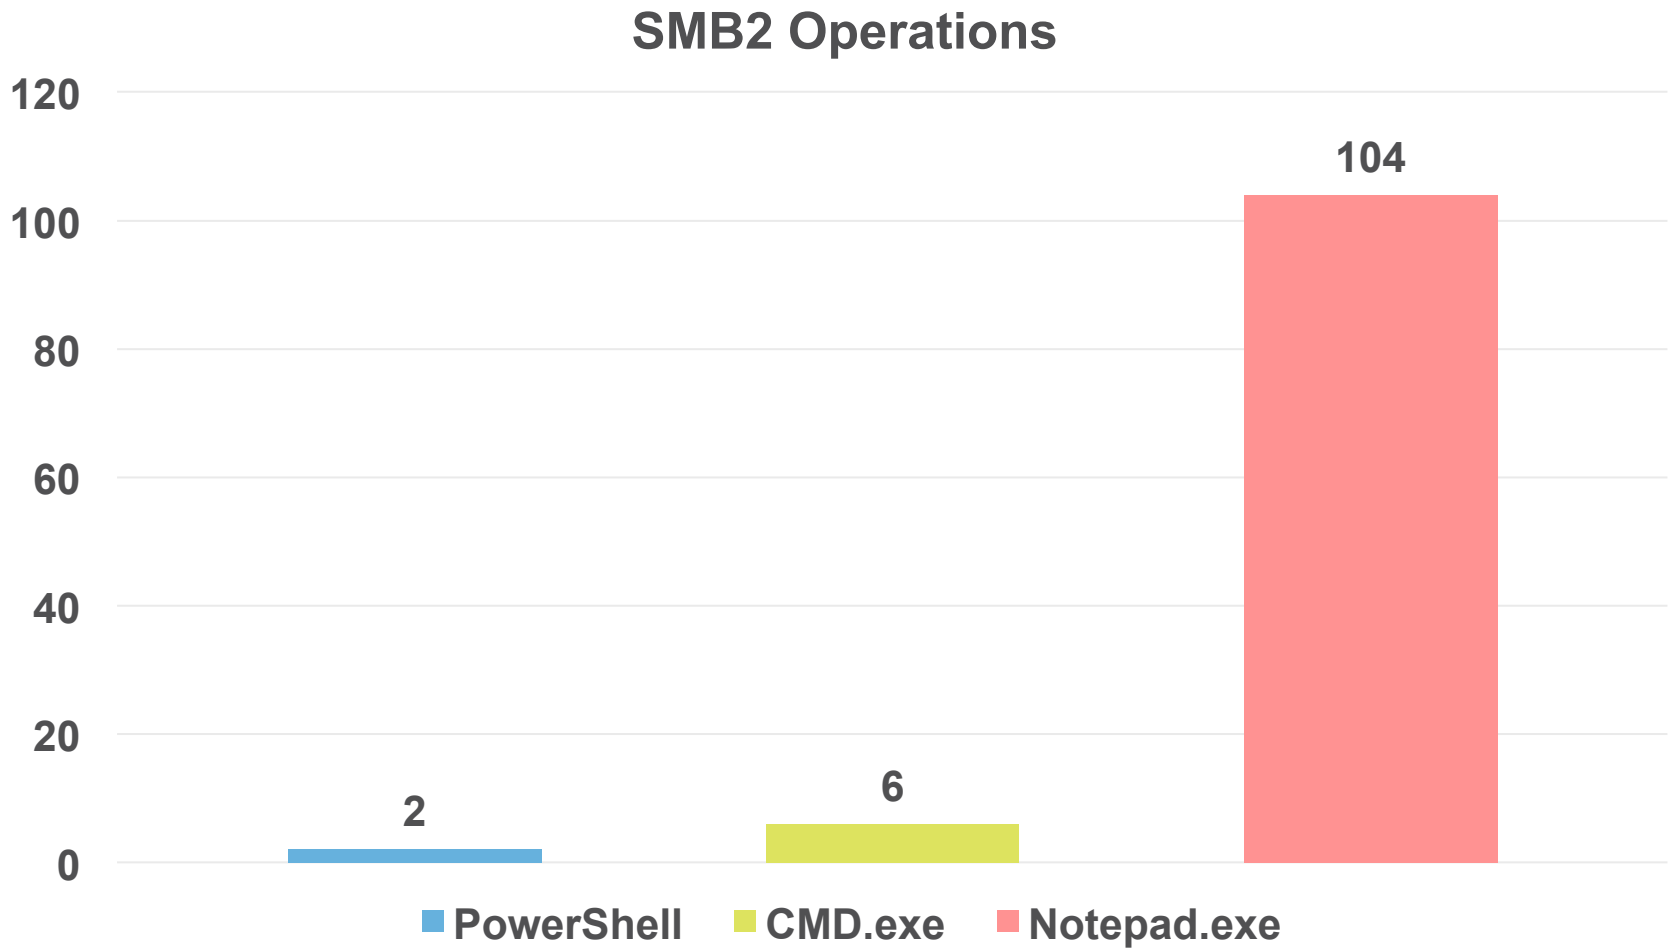

In The File World…

THERE ARE MORE OPS THAN JUST

READS AND WRITES

Agenda

INTRO	INTRODUCTION
FILESYSTEM	FILESYSTEM LAYOUT FUNDAMENTALS
FS OPS	WHAT ARE FILESYSTEM OPERATIONS?
FILE PERF	**PERFORMANCE IN A FILE ORIENTED WORLD**
FUN	GRAPH FUN WITH PERFORMANCE
END	SUMMARY

Performance In a File Oriented World is Affected By …

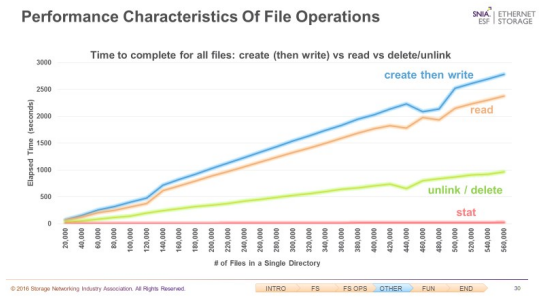

Performance Differences Between Operations

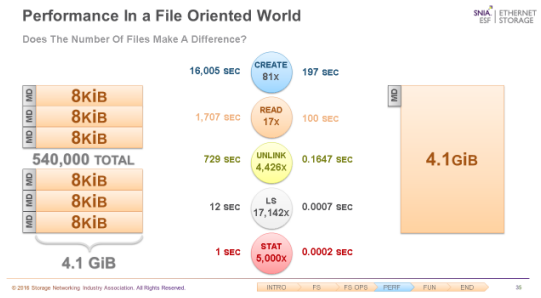

Data vs Metadata Ratios

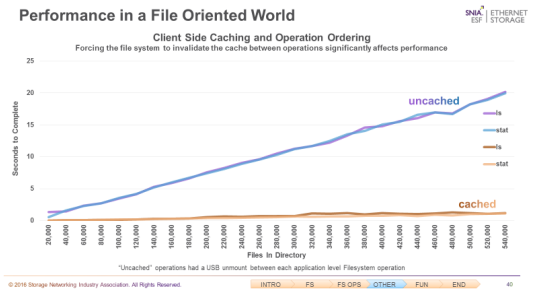

Caching

Performance Characteristics Of File Operations

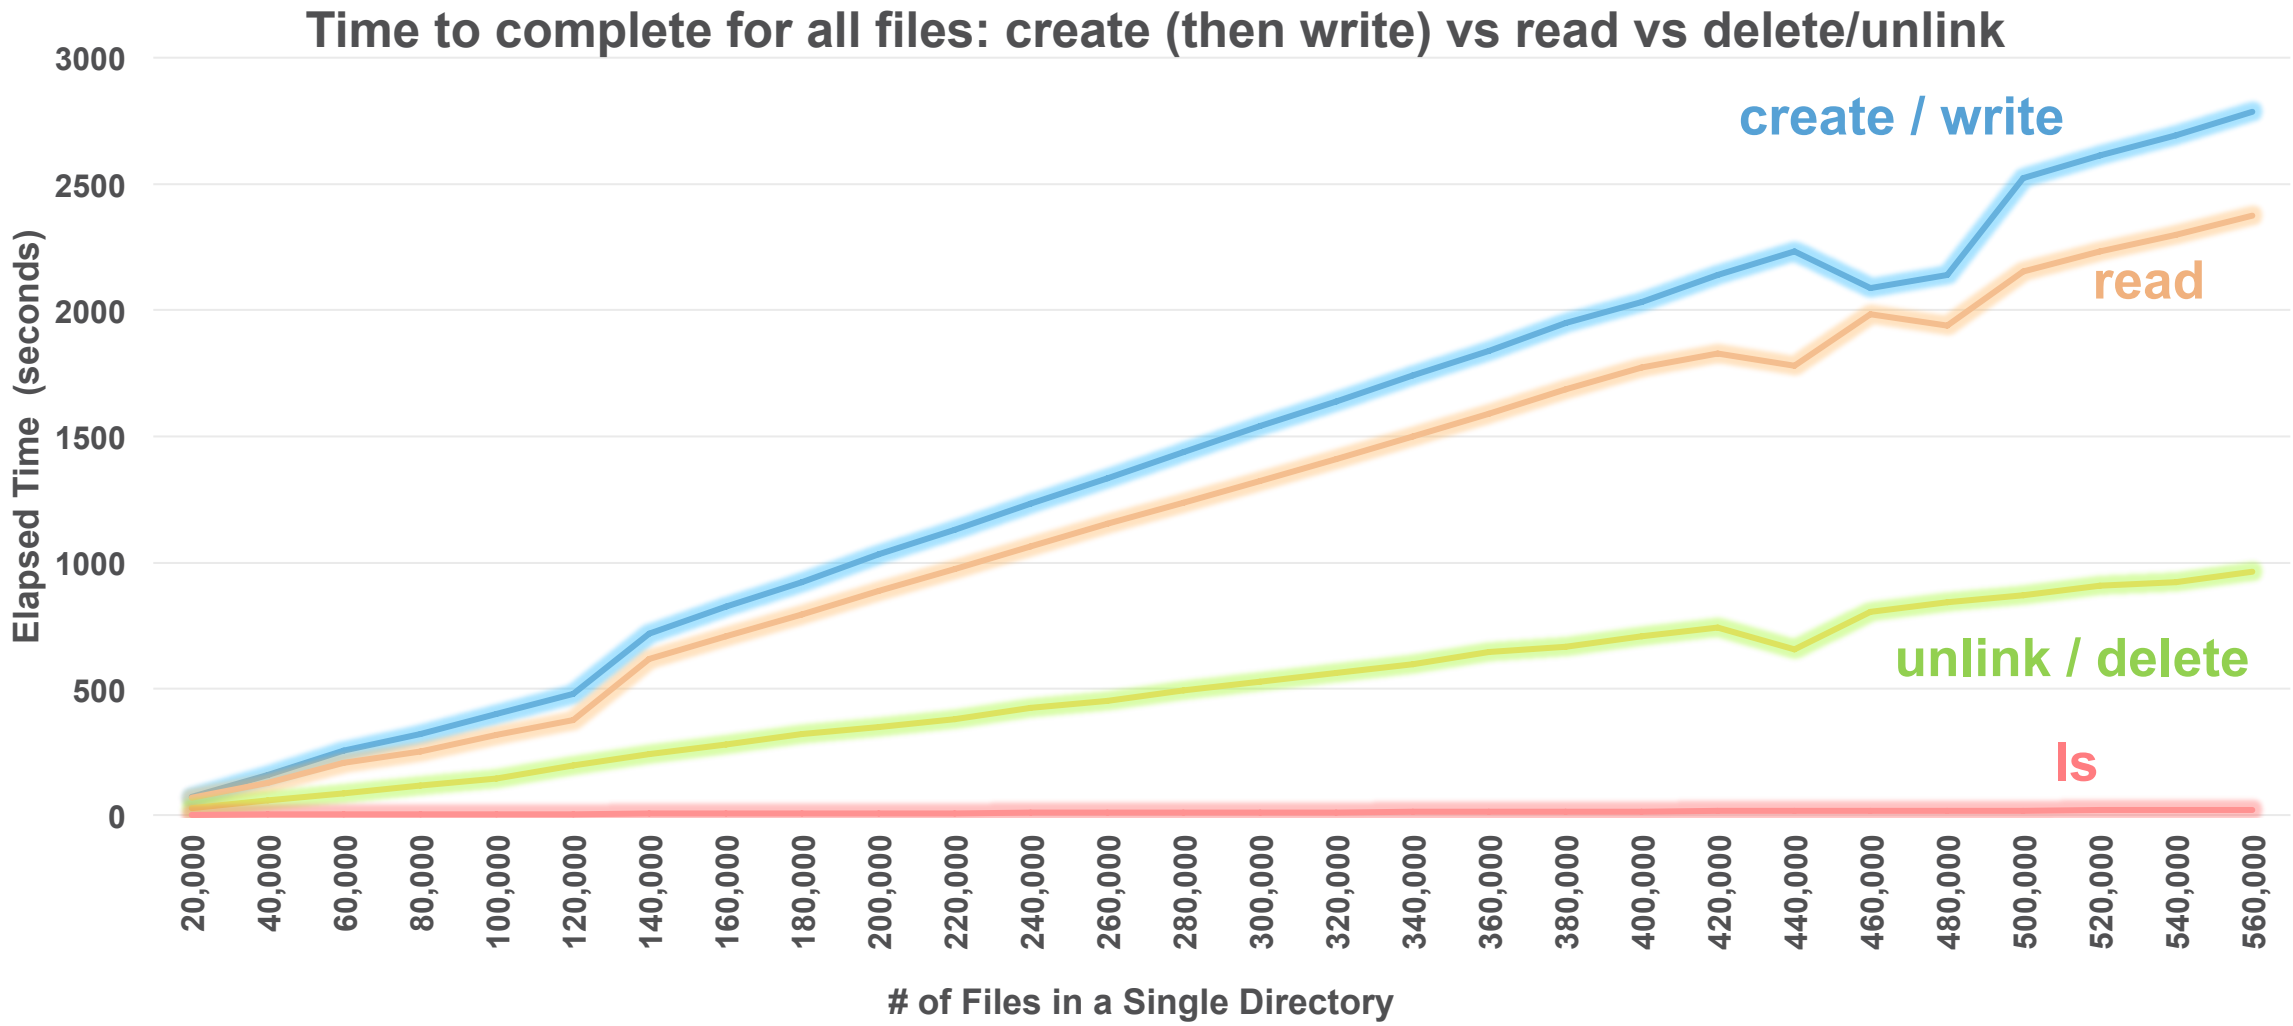

Performance Characteristics Of File Operations

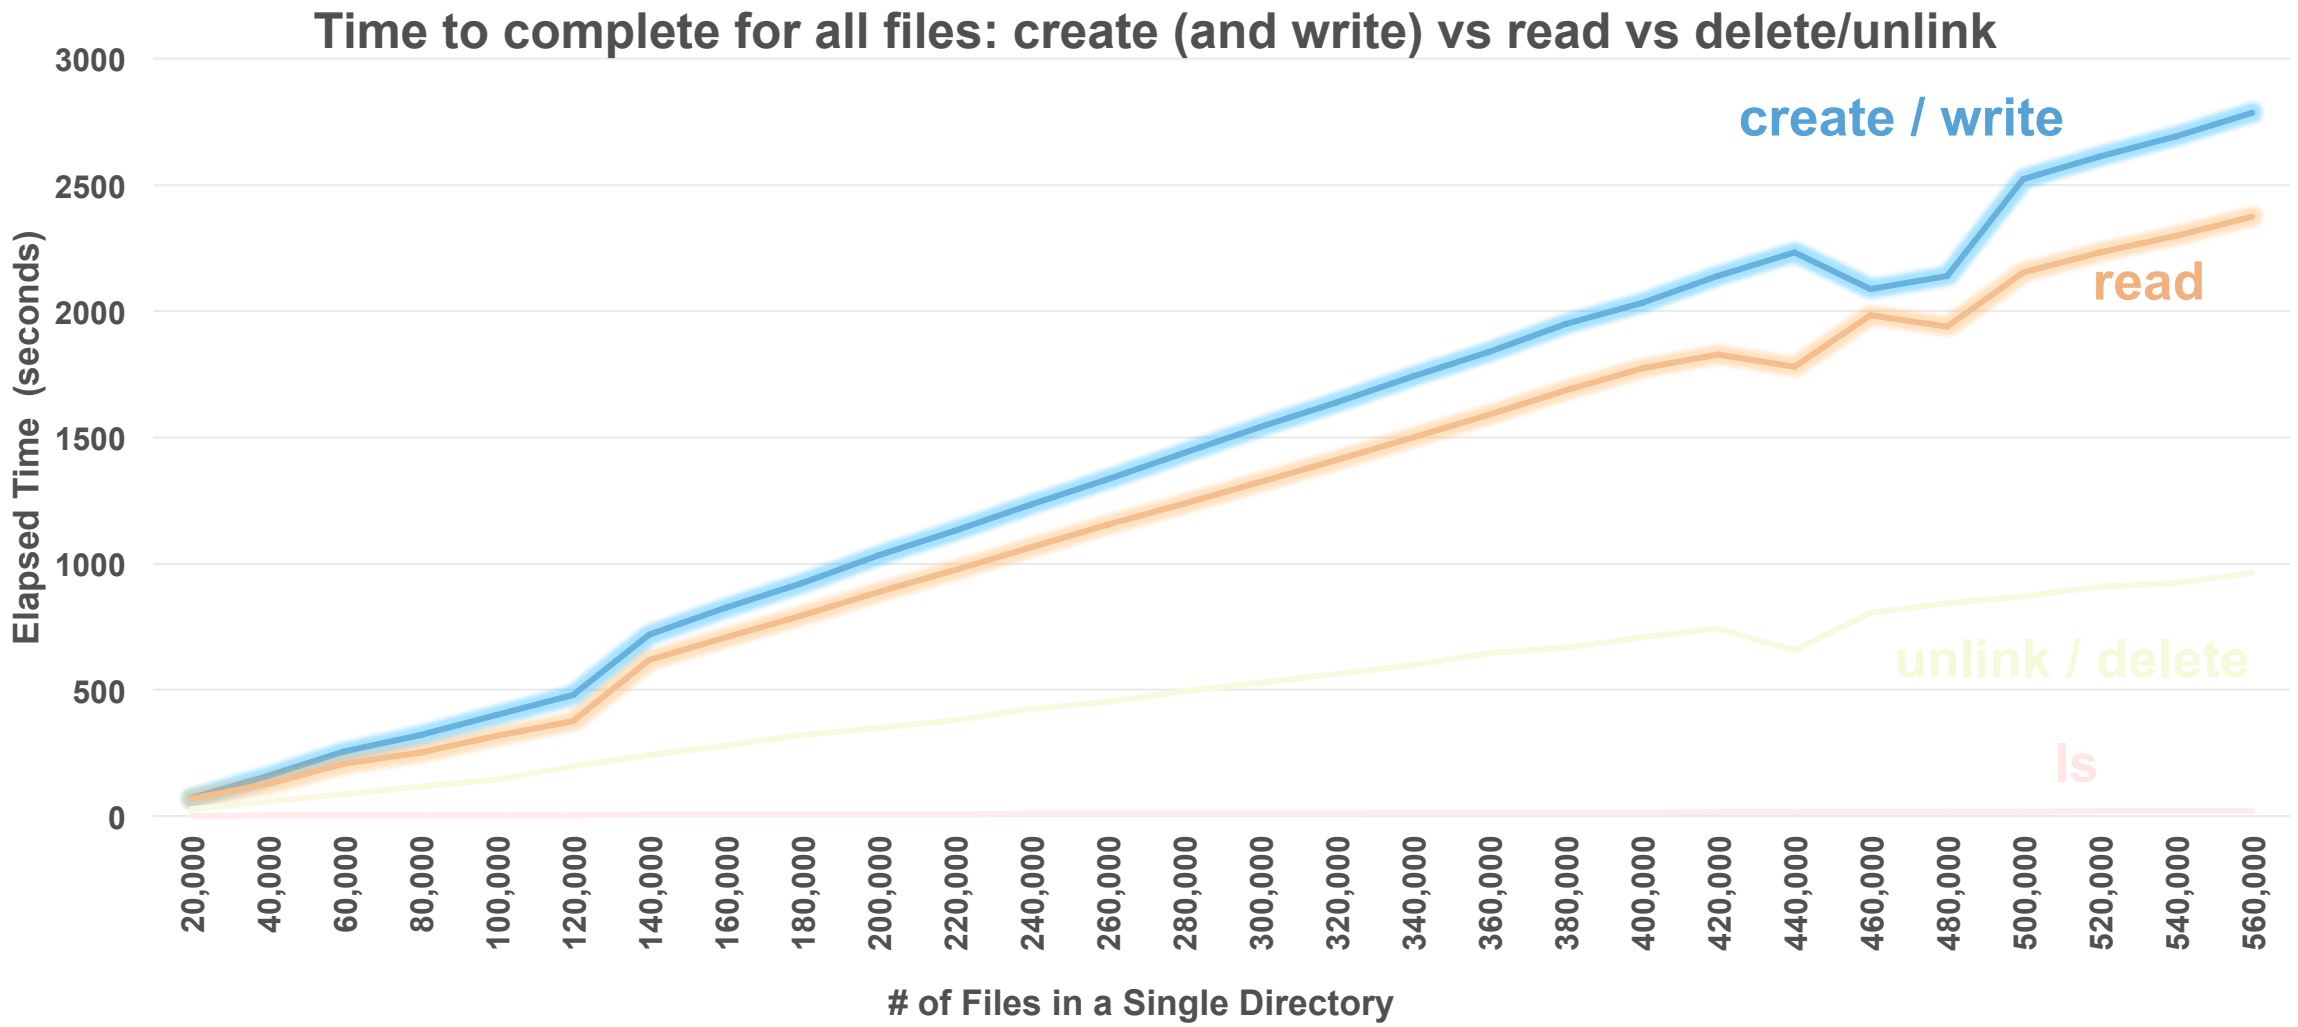

Performance Characteristics Of File Operations

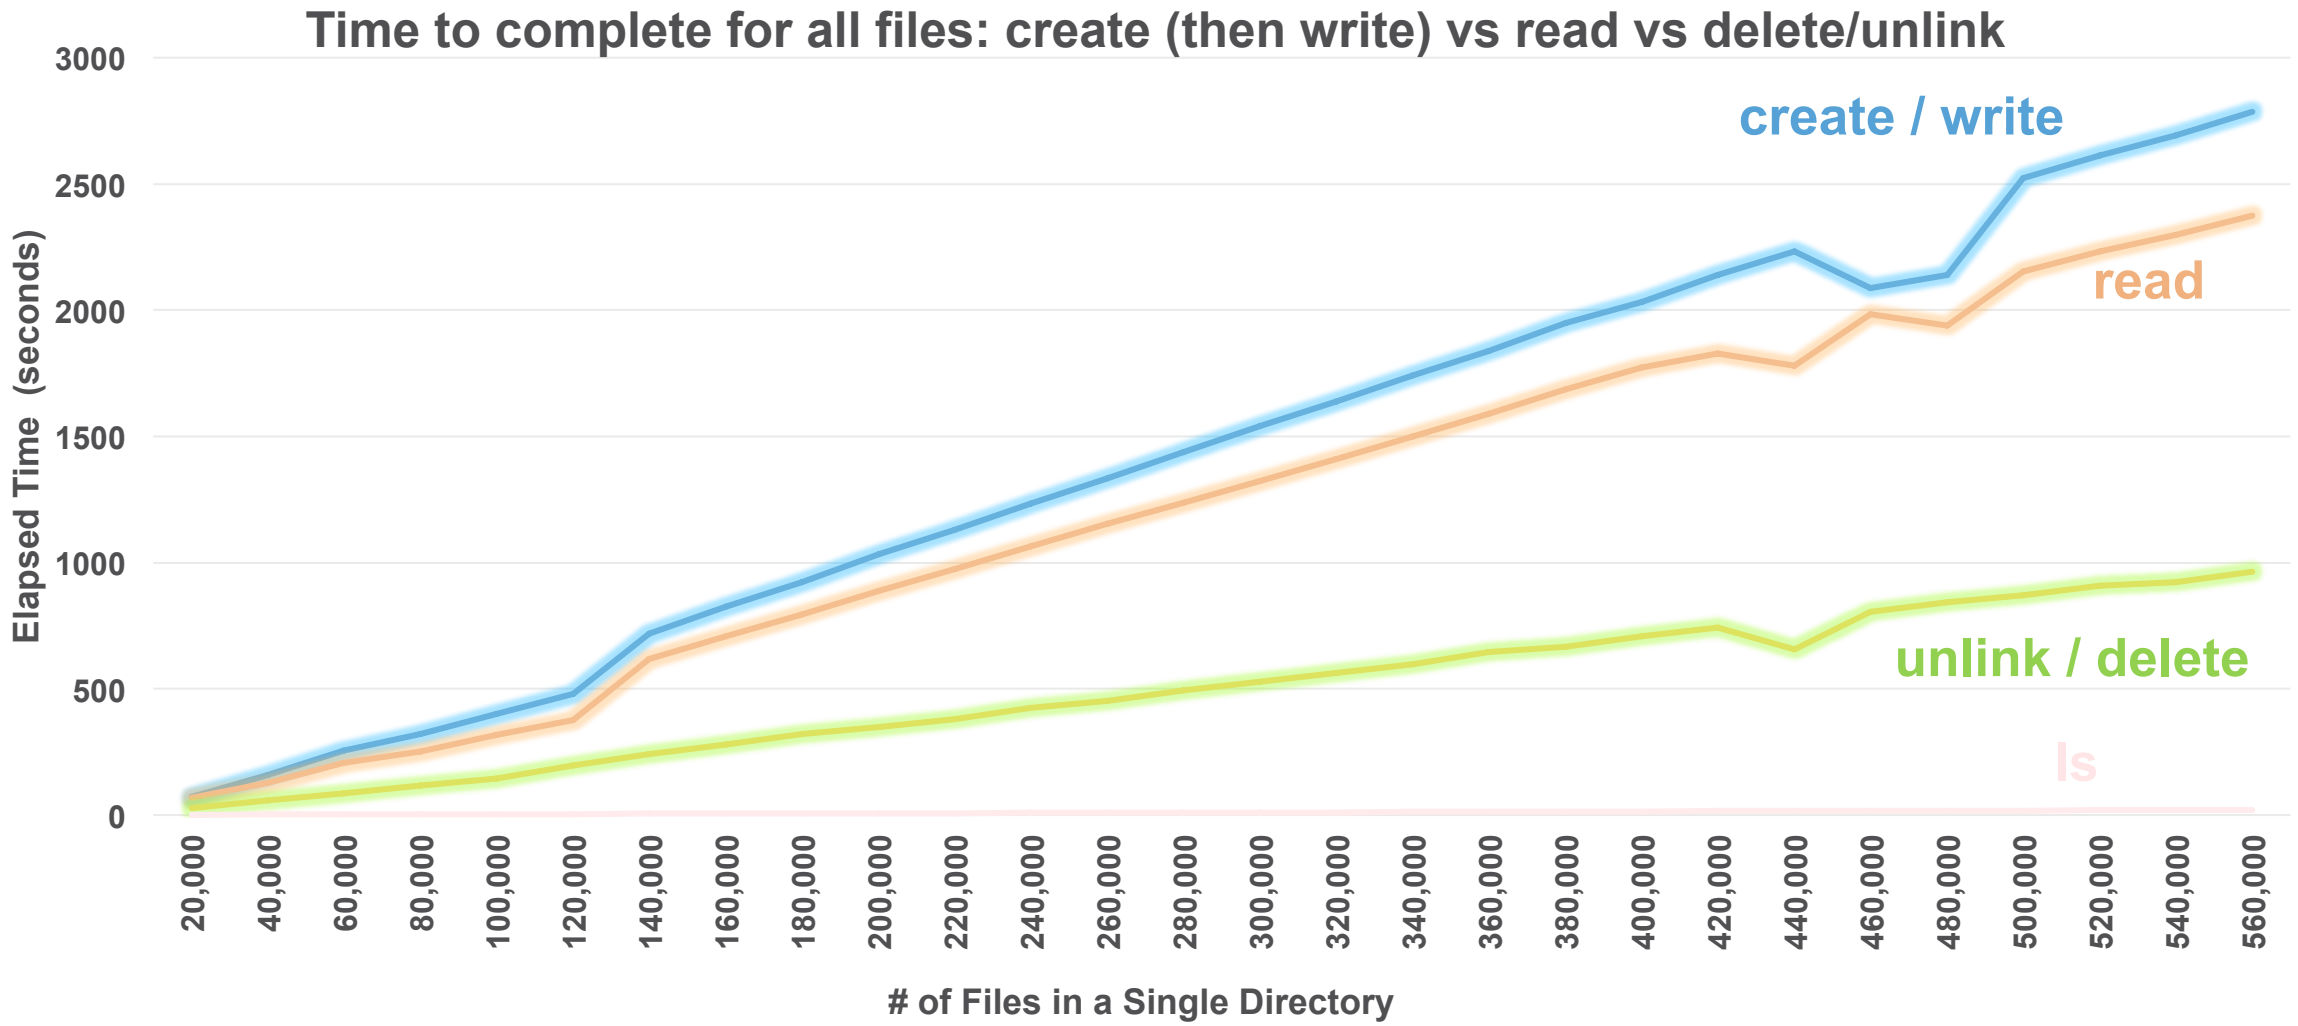

Performance Characteristics Of File Operations

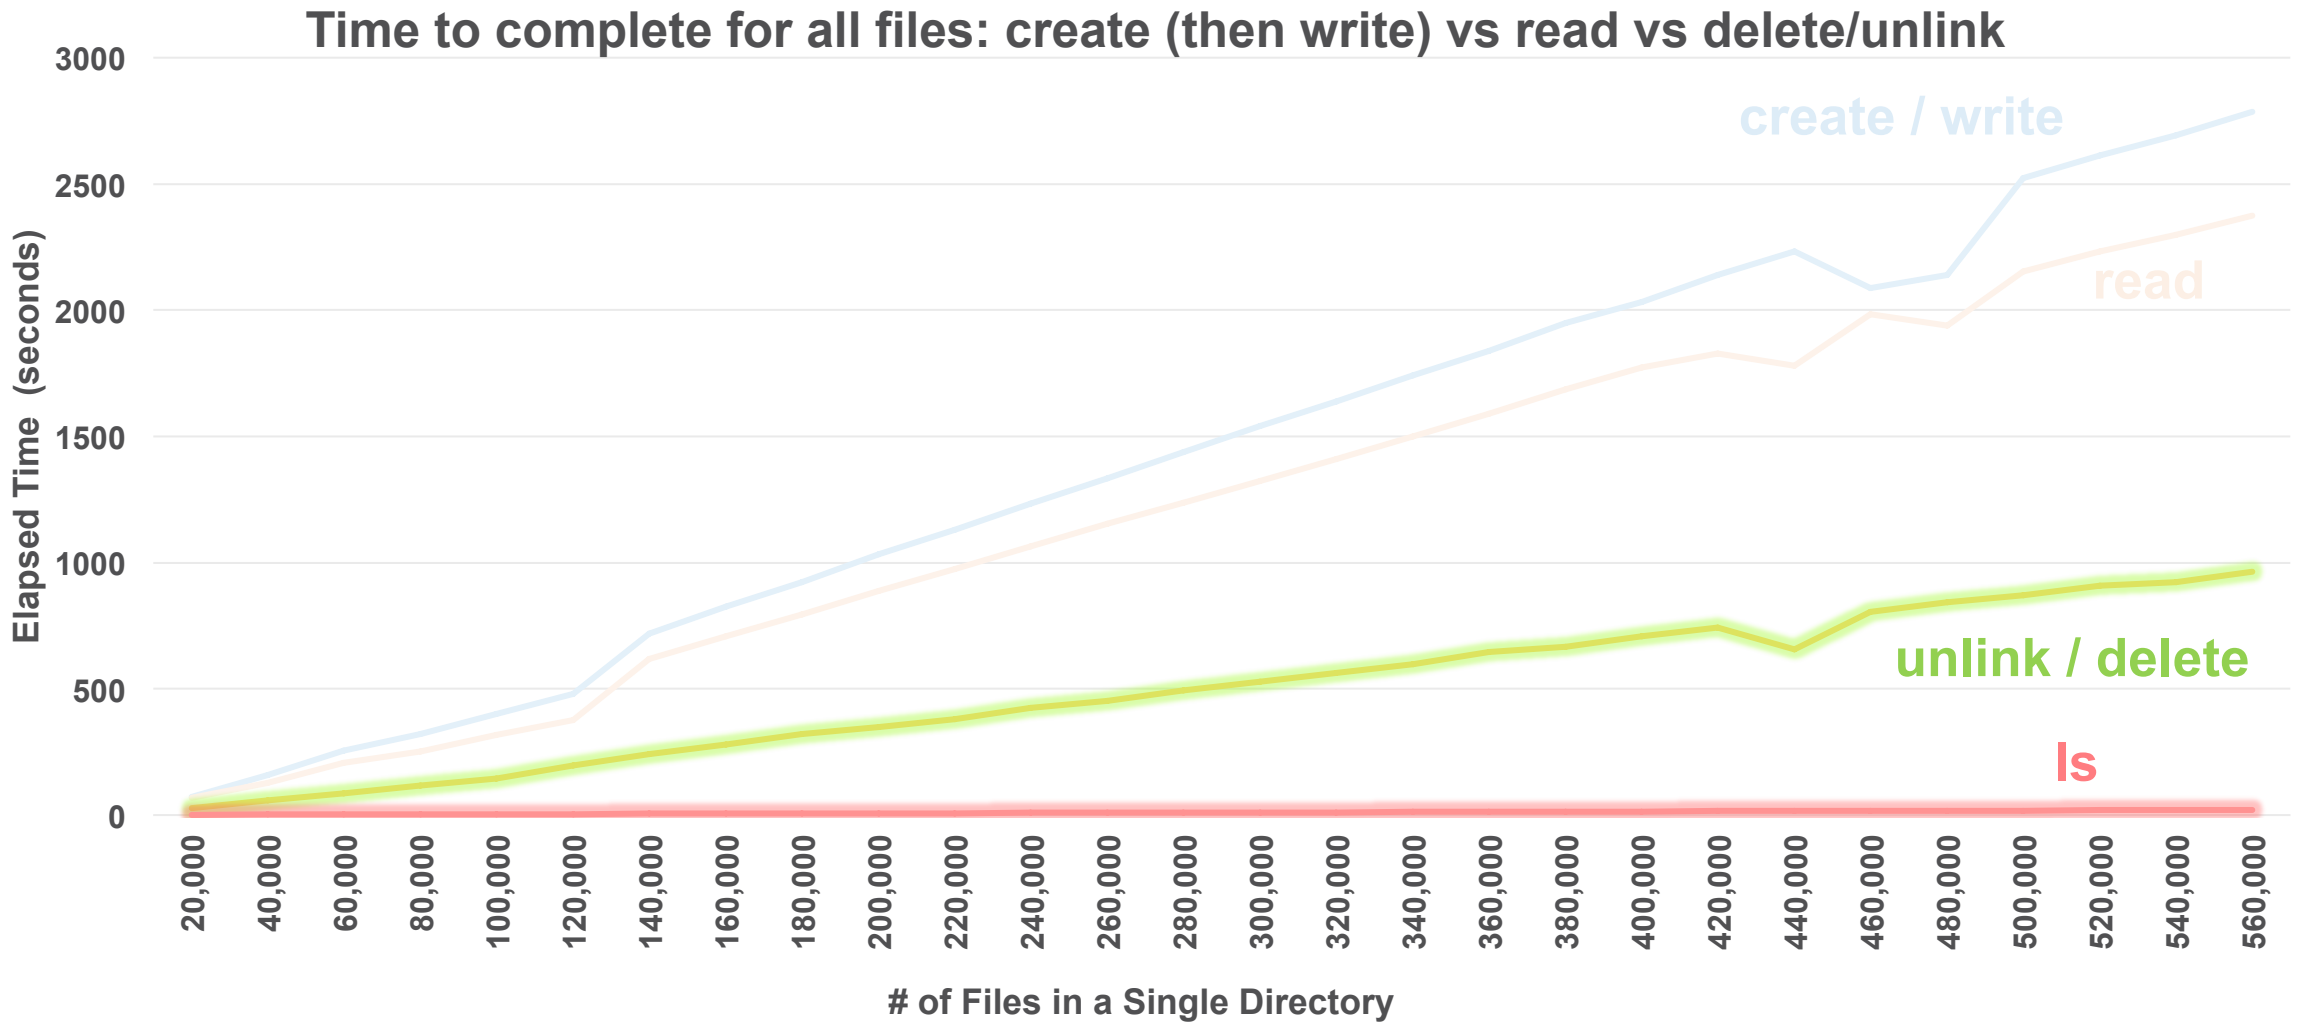

Performance In a File Oriented World

SNIA ESF | ETHERNET STORAGE

Be Careful … Operation Performance Will Almost Always Be Affected By the SUT

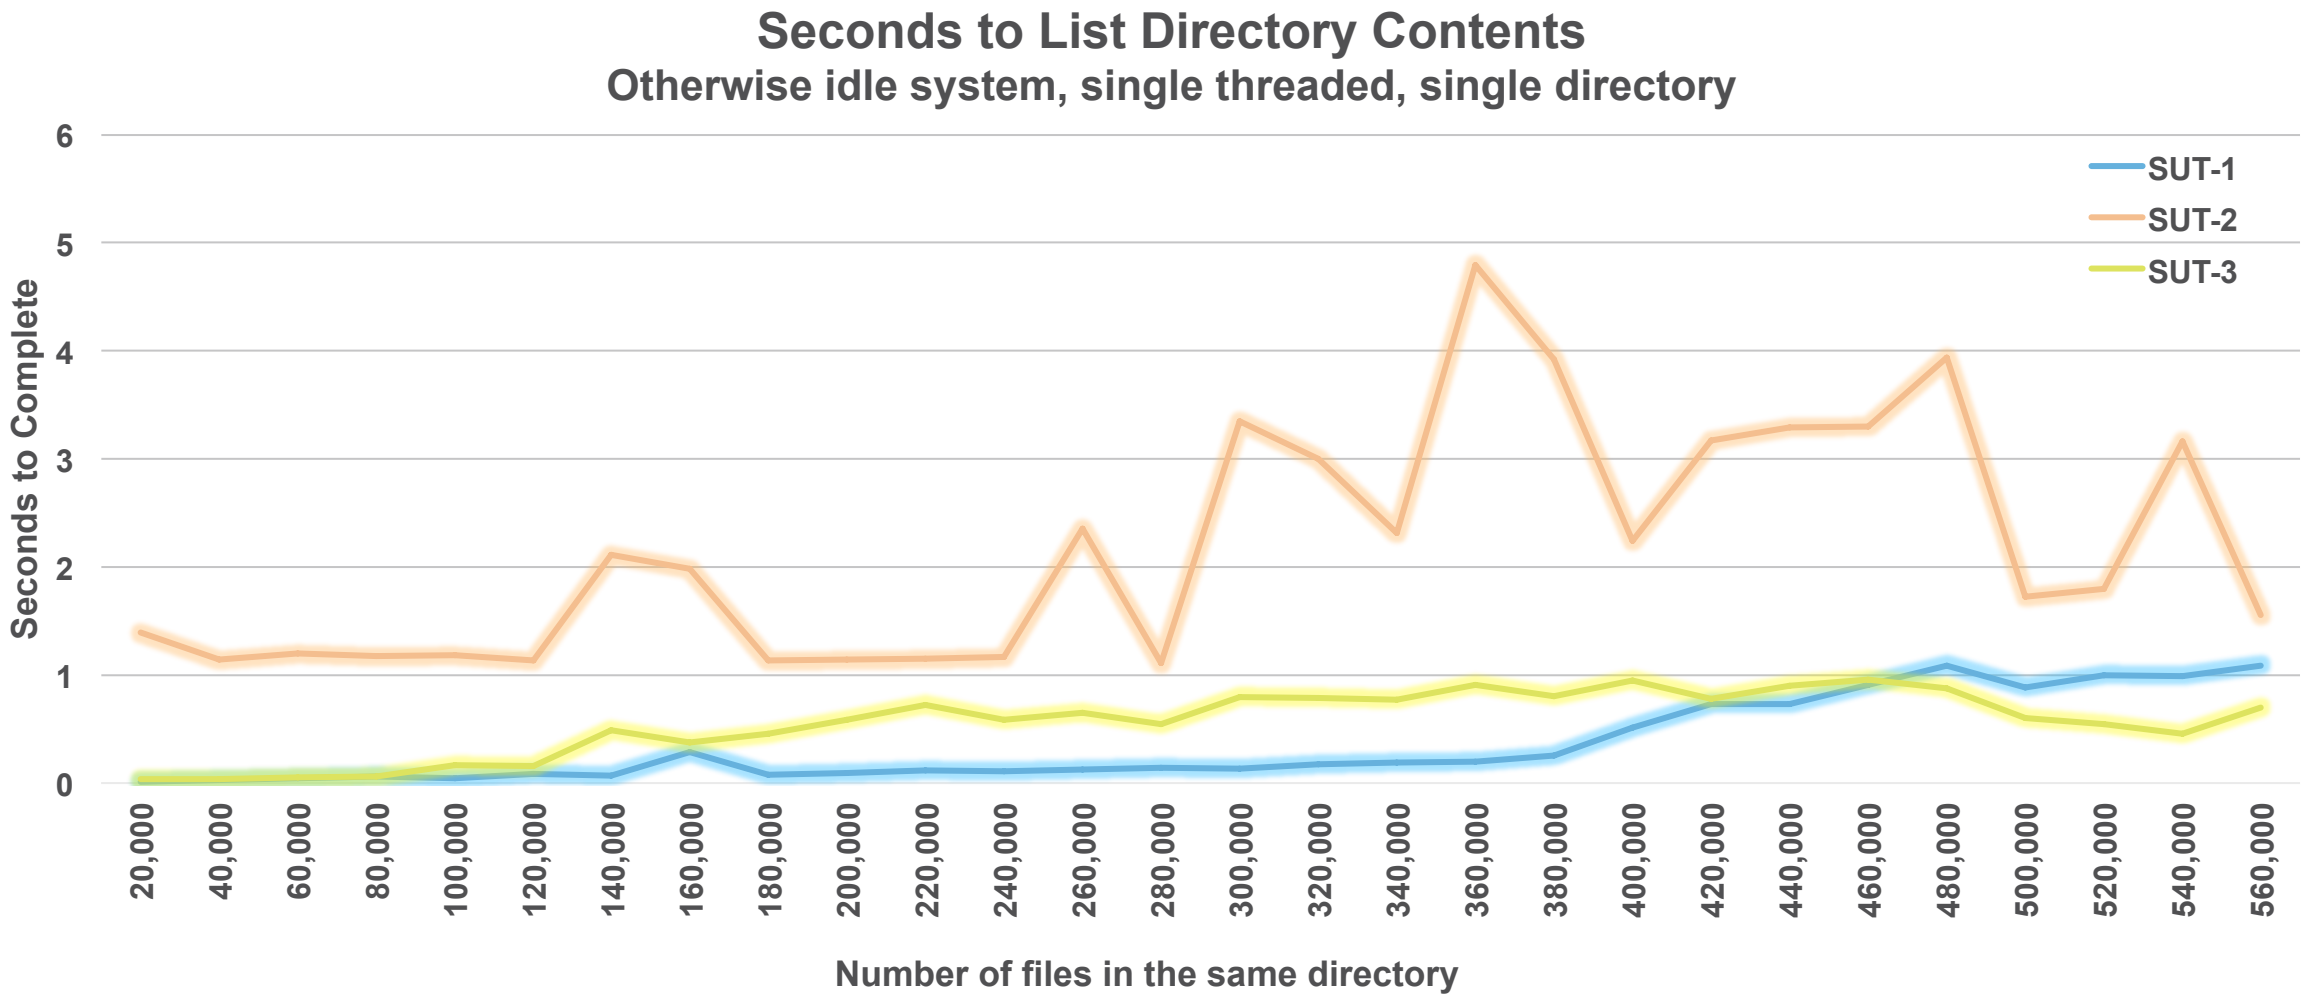

Performance In a File Oriented World is Affected By …

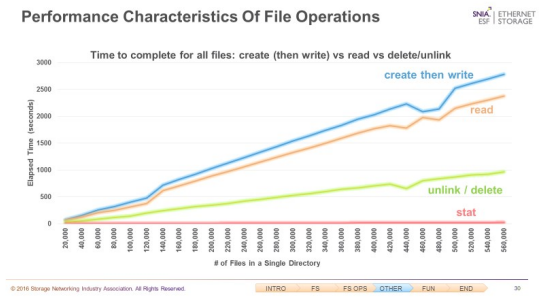

Performance Differences Between Operations

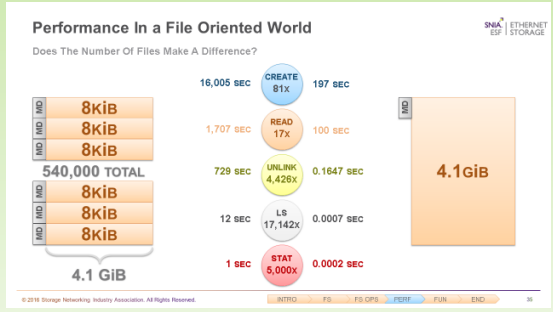

Data vs Metadata Ratios

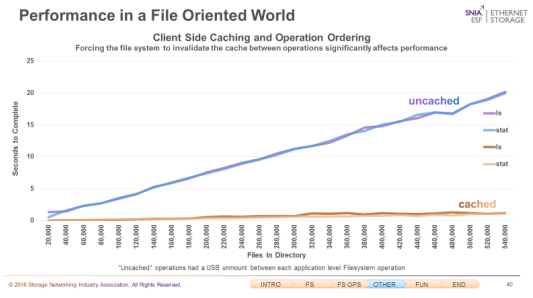

Caching

Performance In a File Oriented World

Does The Number Of Files Make A Difference?

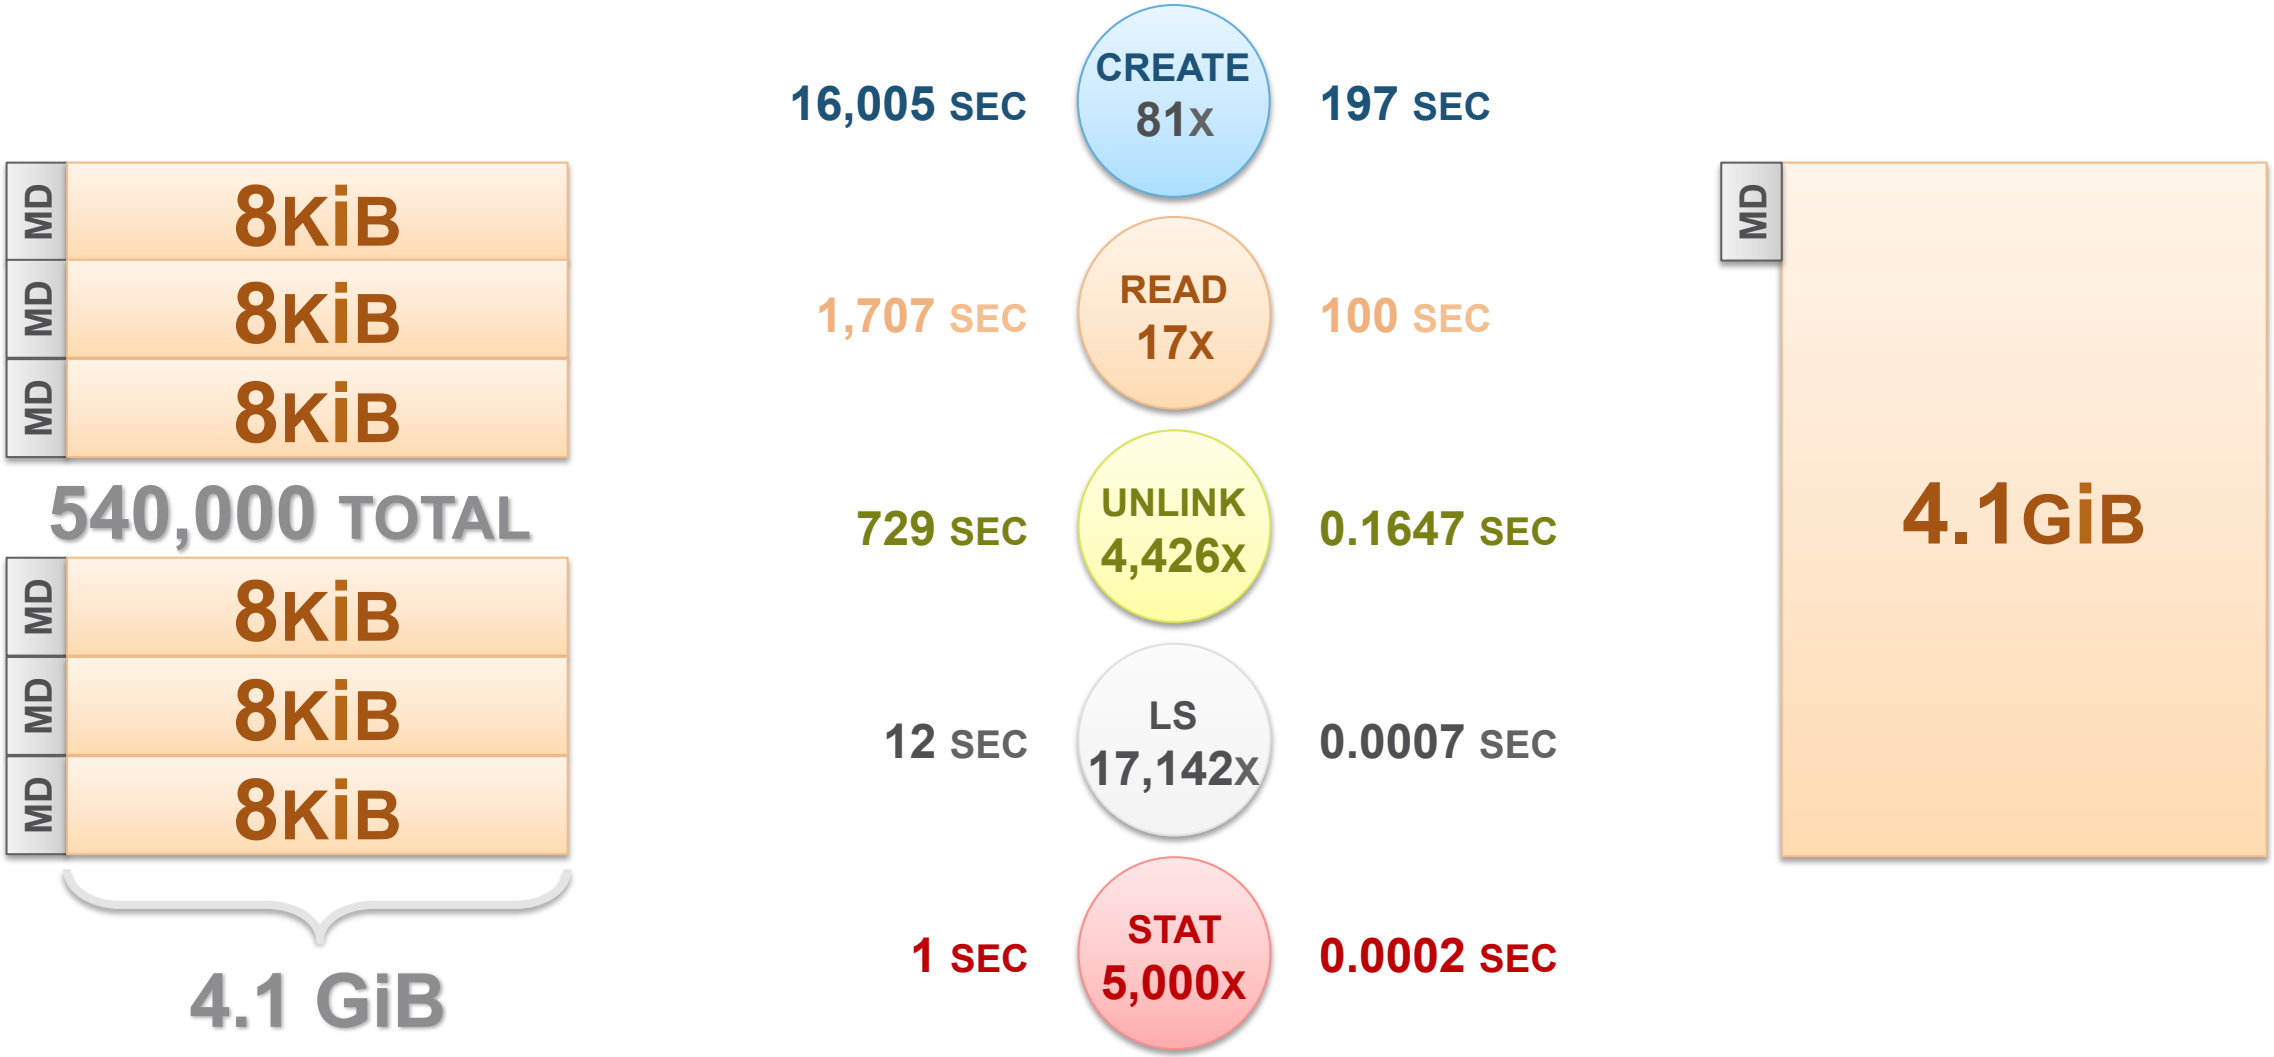

Performance In a File Oriented World is Affected By …

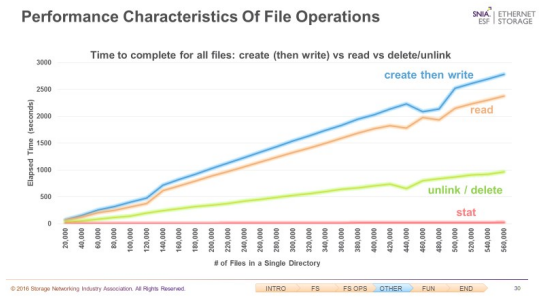

Performance Differences Between Operations

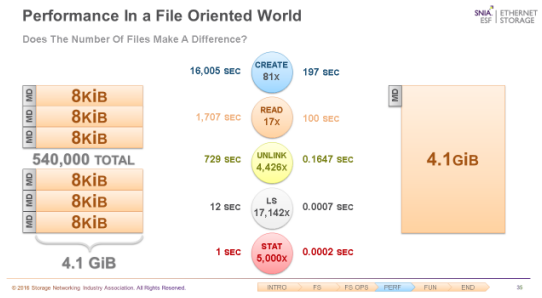

Data vs Metadata Ratios

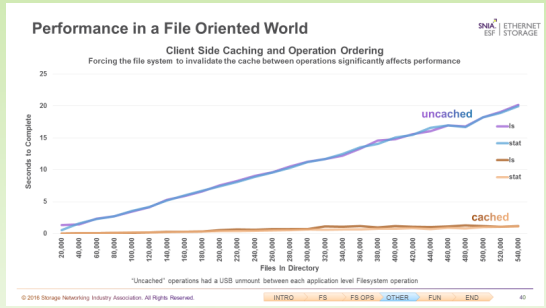

Caching

Performance in a File Oriented World

Client Side Caching and Operation Ordering

Forcing the file system to invalidate the cache between operations significantly affects performance

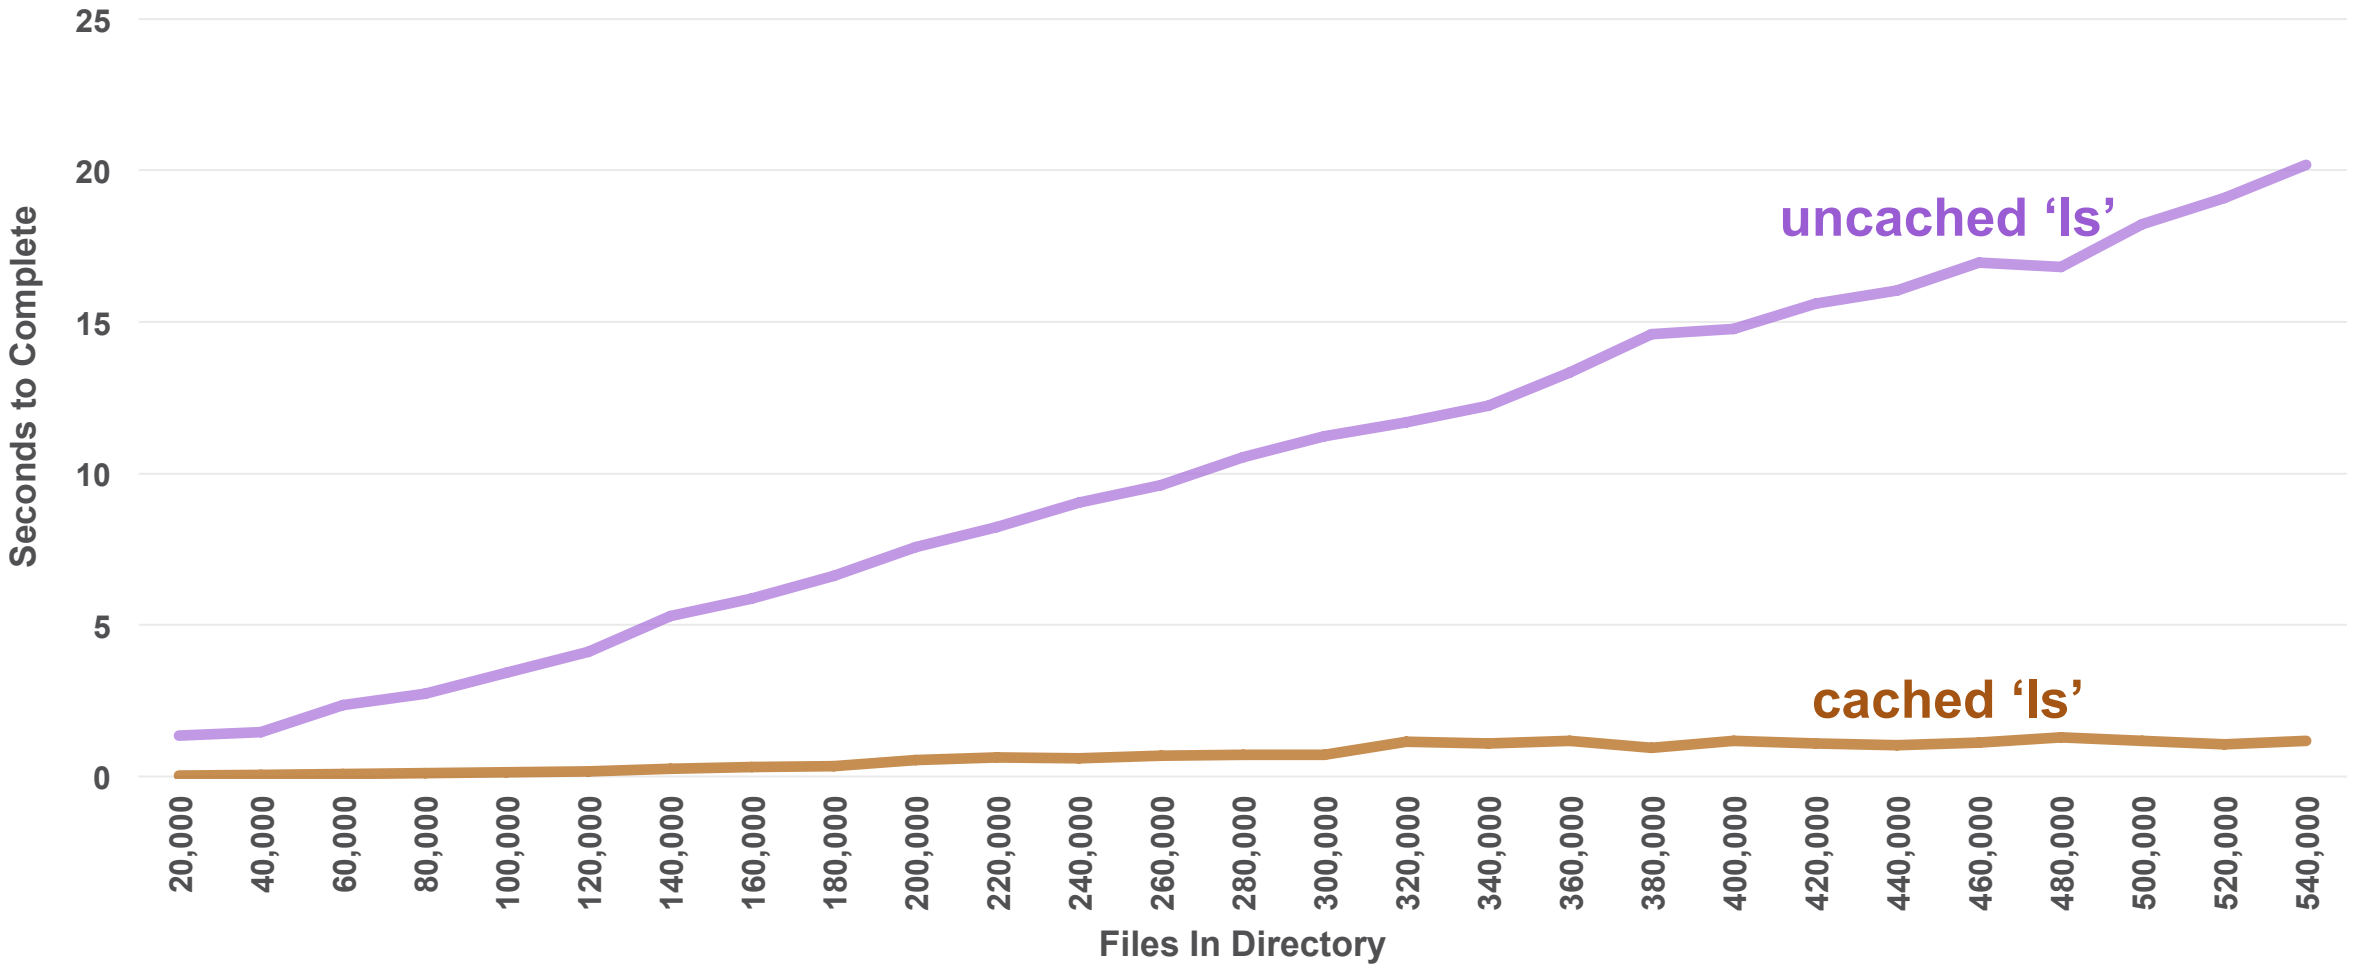

"Uncached 'ls'" had a USB unmount just prior to the 'ls' command execution

Performance In a File Oriented World is Affected By …

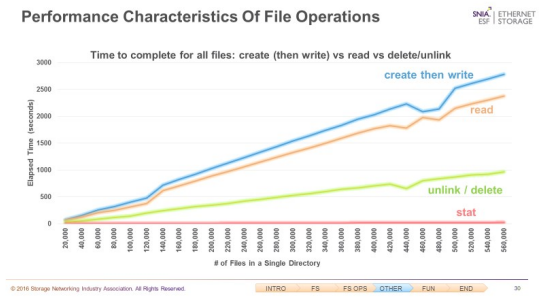

Performance Differences Between Operations

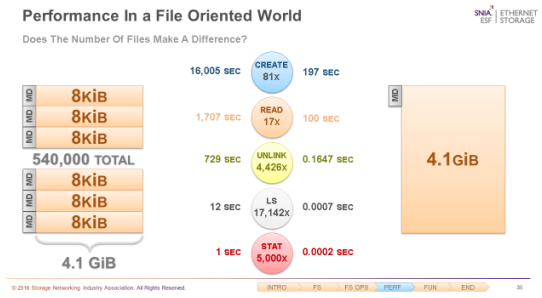

Data vs Metadata Ratios

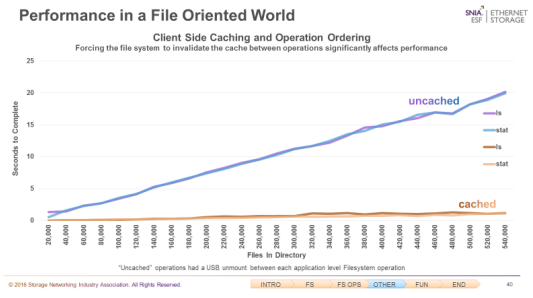

Caching

Agenda

INTRO	INTRODUCTION
FILESYSTEM	FILESYSTEM LAYOUT FUNDAMENTALS
FS OPS	WHAT ARE FILESYSTEM OPERATIONS?
FILE PERF	PERFORMANCE IN A FILE ORIENTED WORLD
FUN	**GRAPH FUN WITH PERFORMANCE**
END	SUMMARY

Which Is Better?

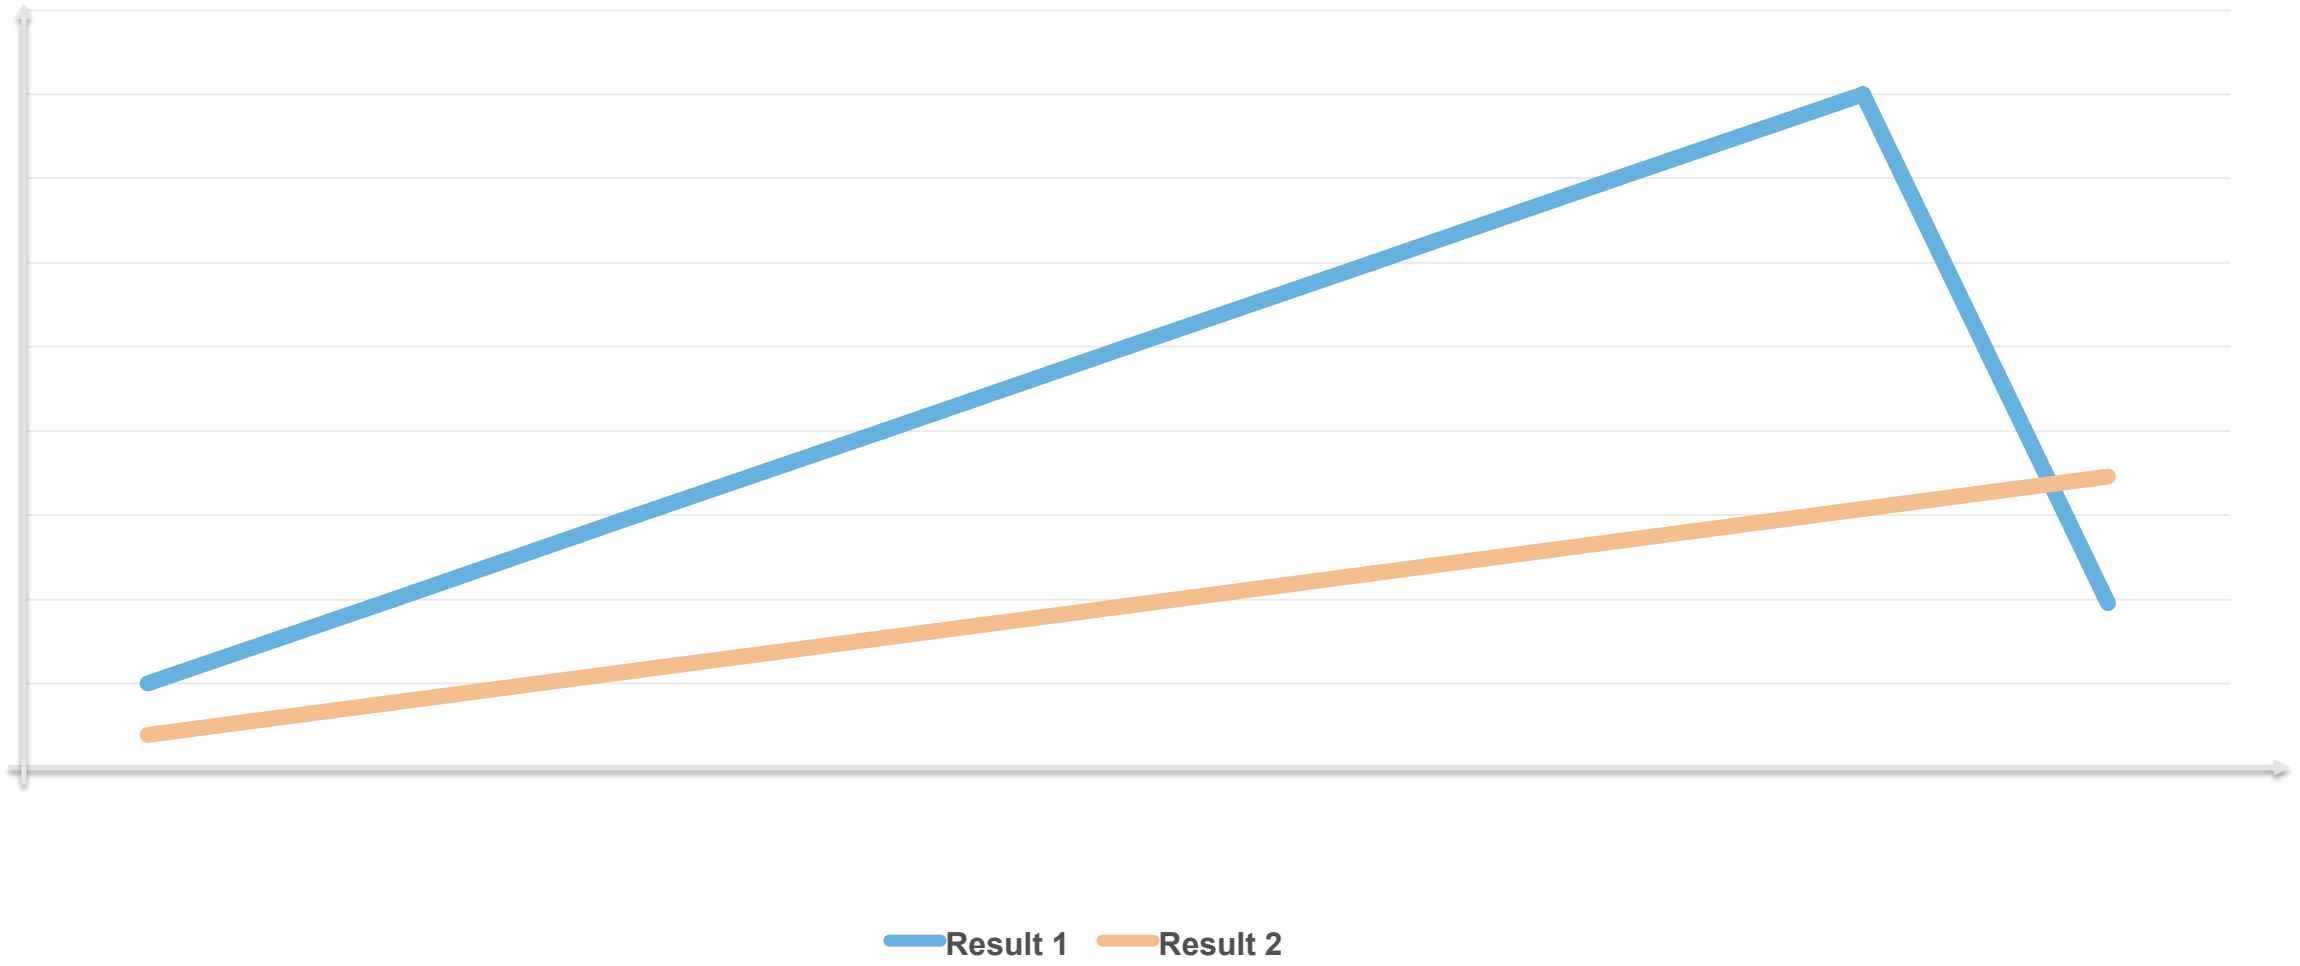

Which Is Better?

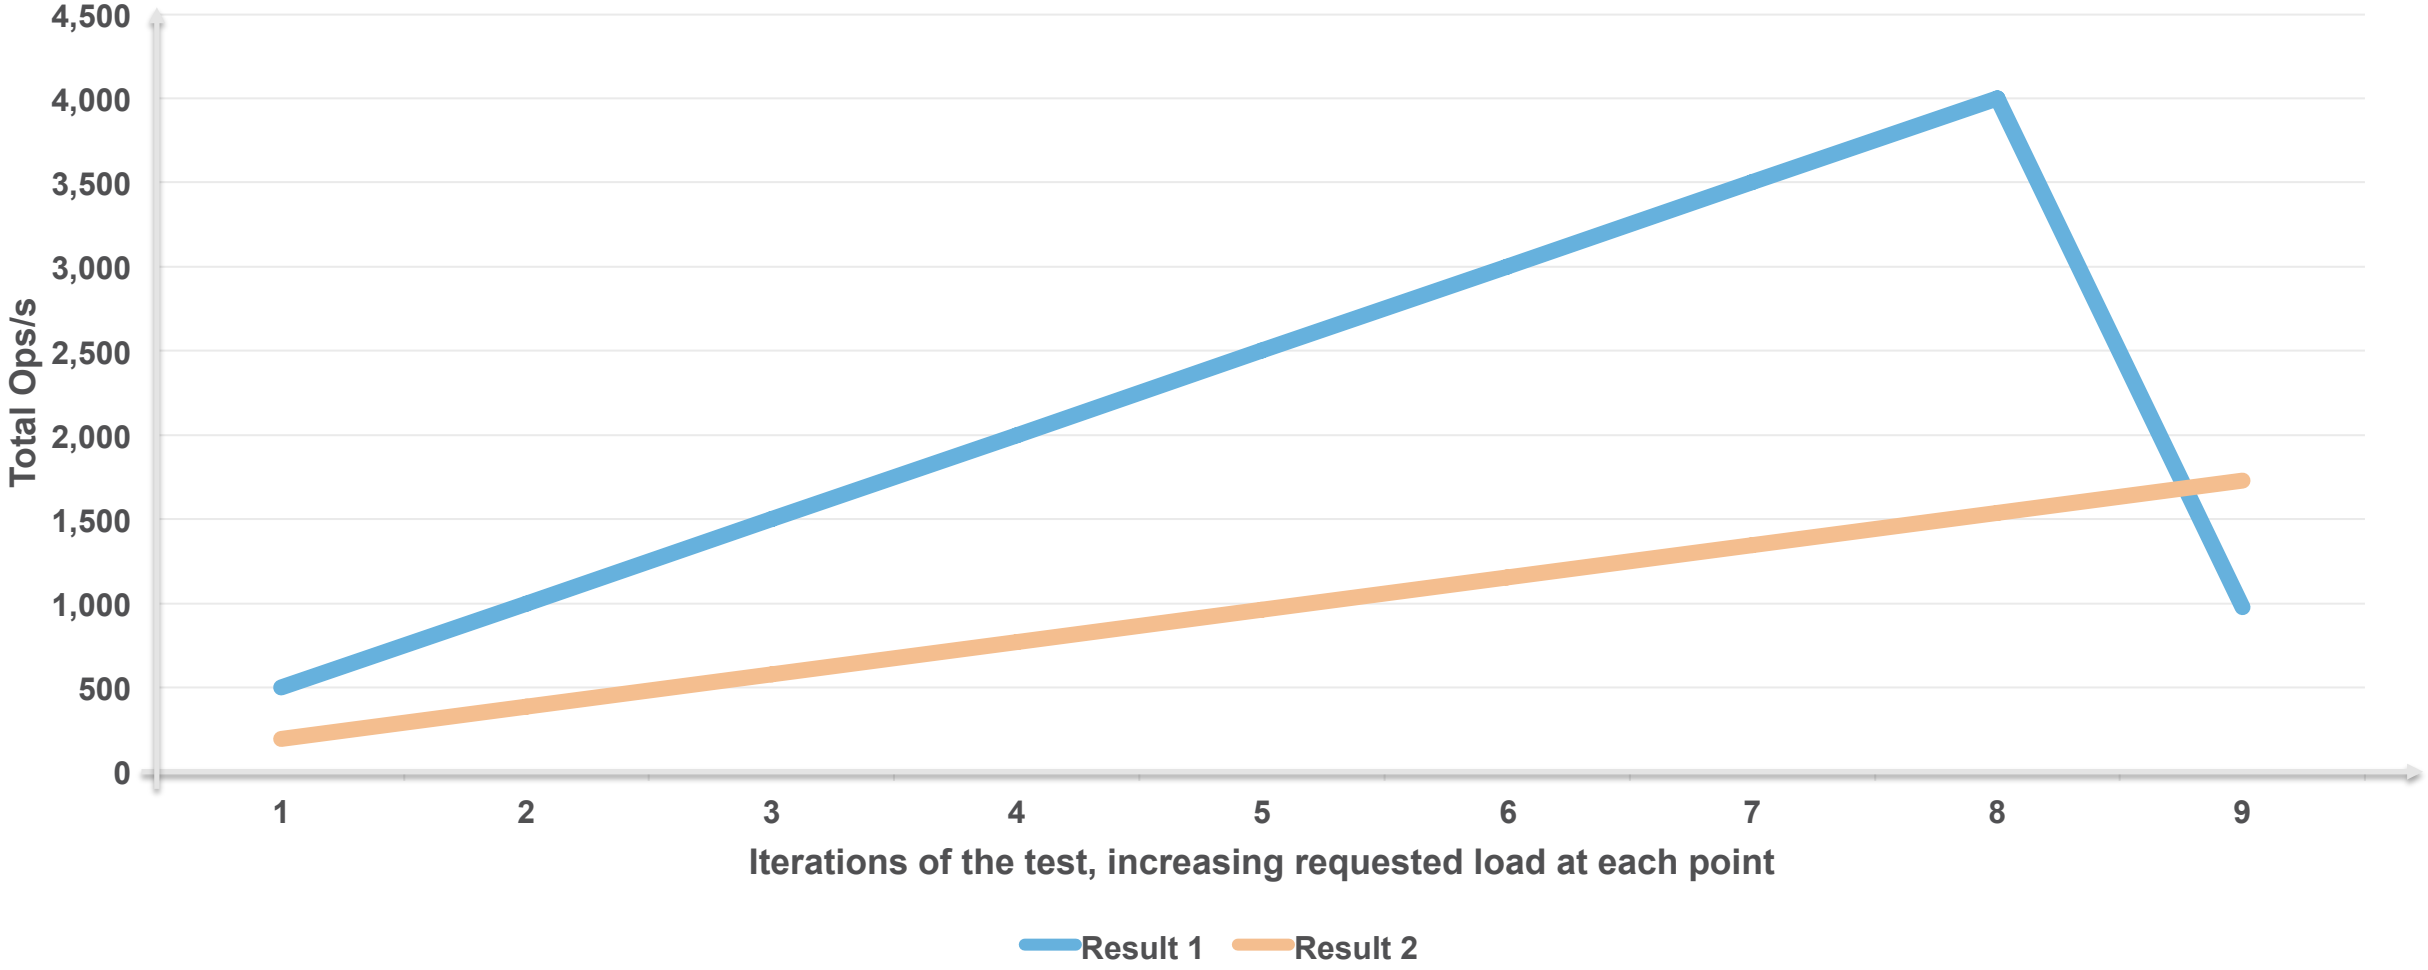

Which Is Better?

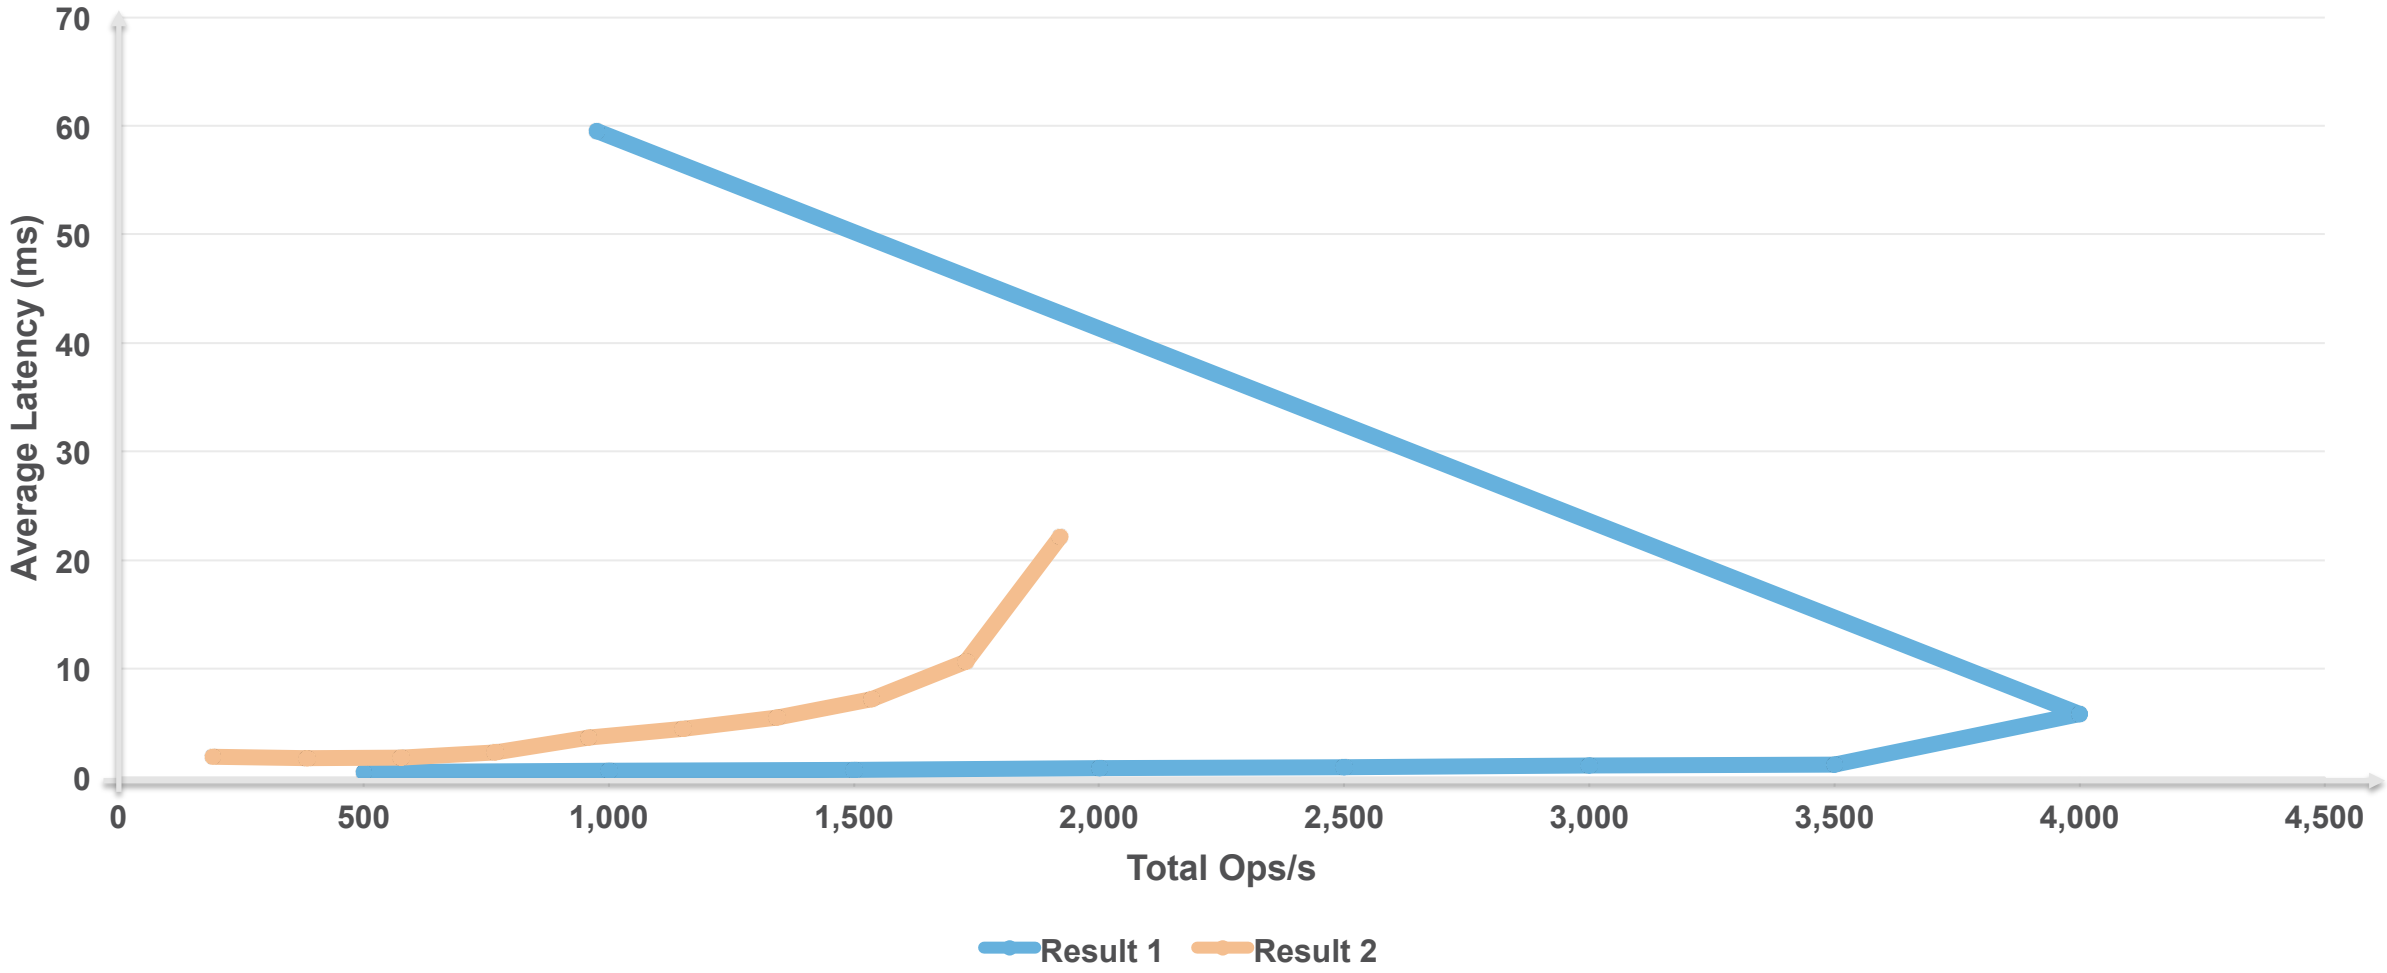

Which Is Better?

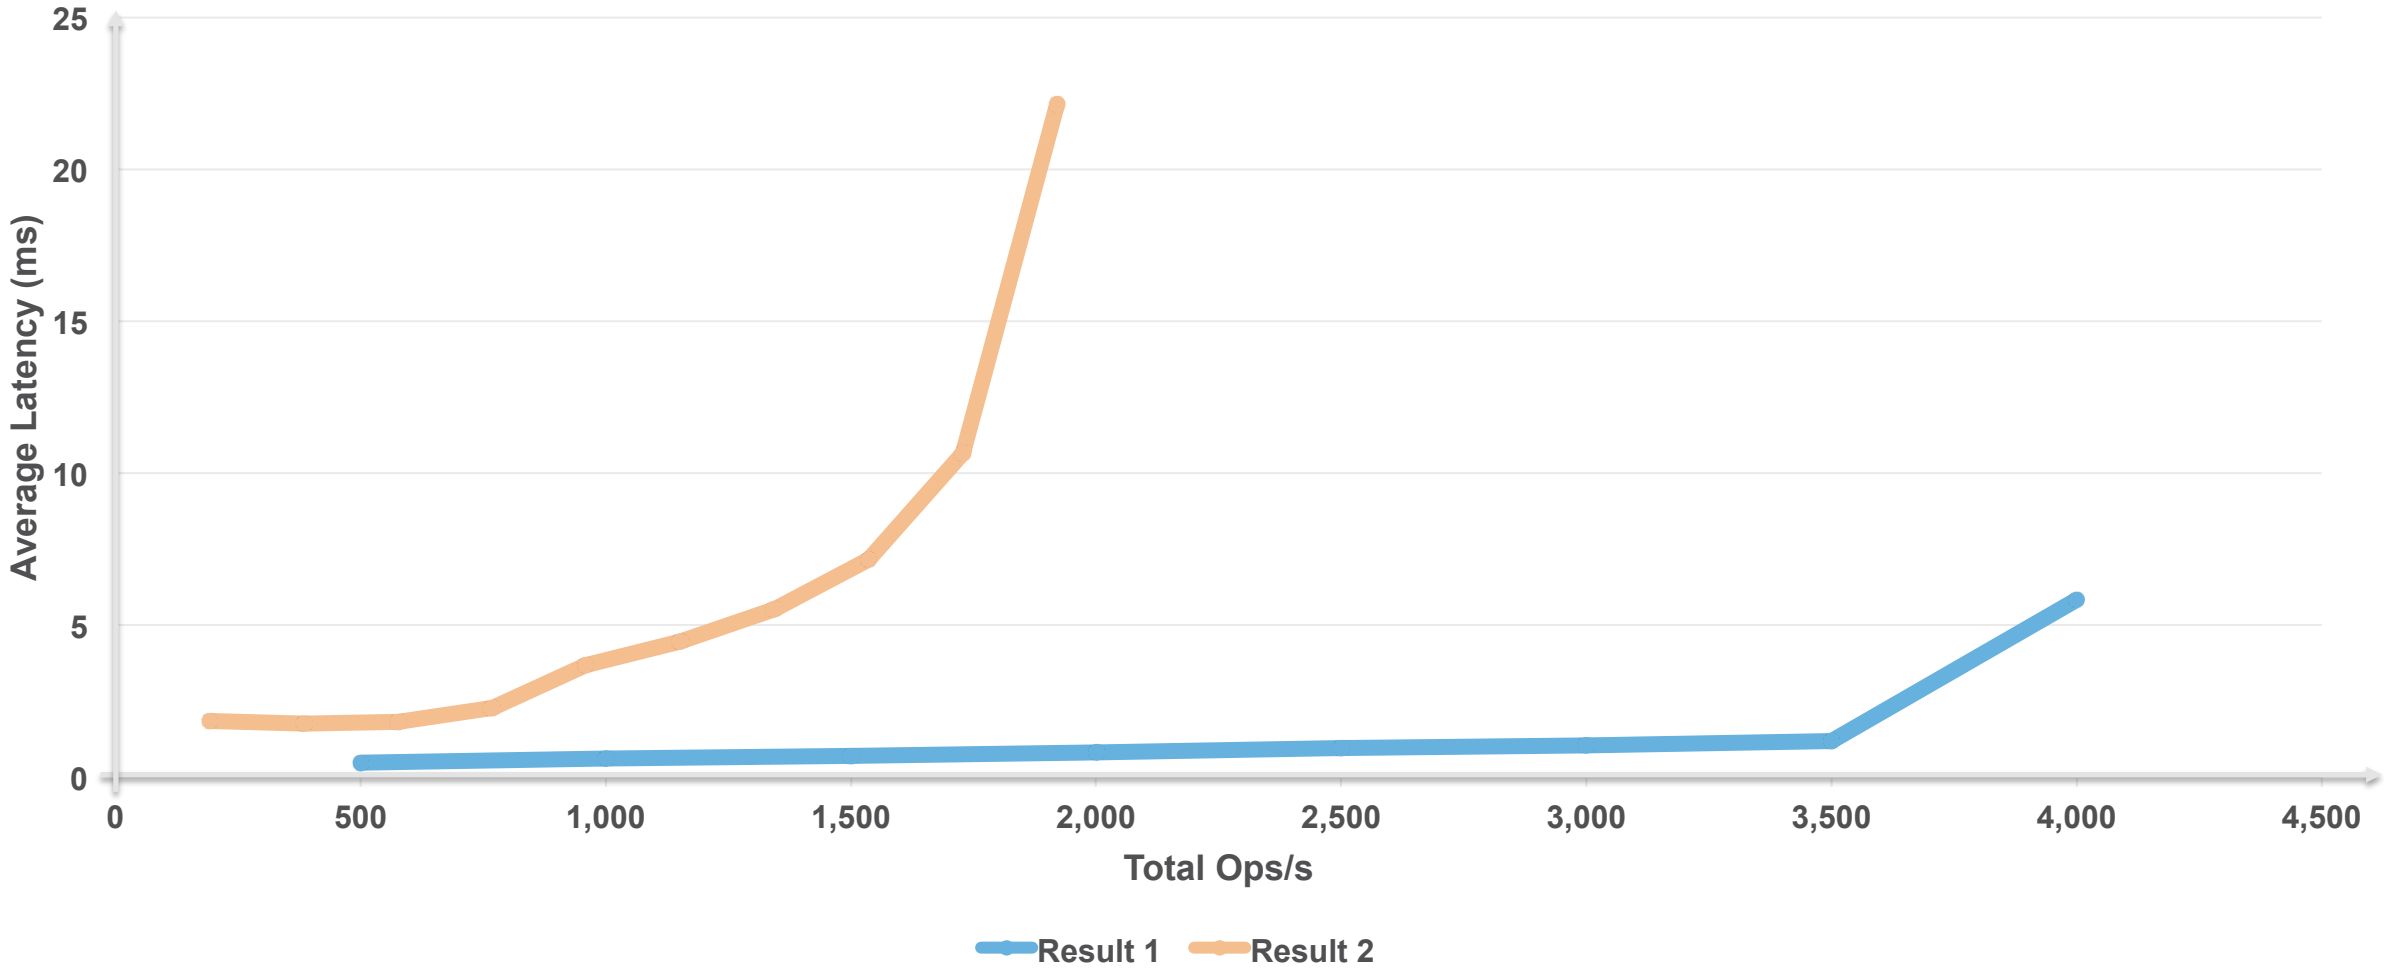

Which Is Better?

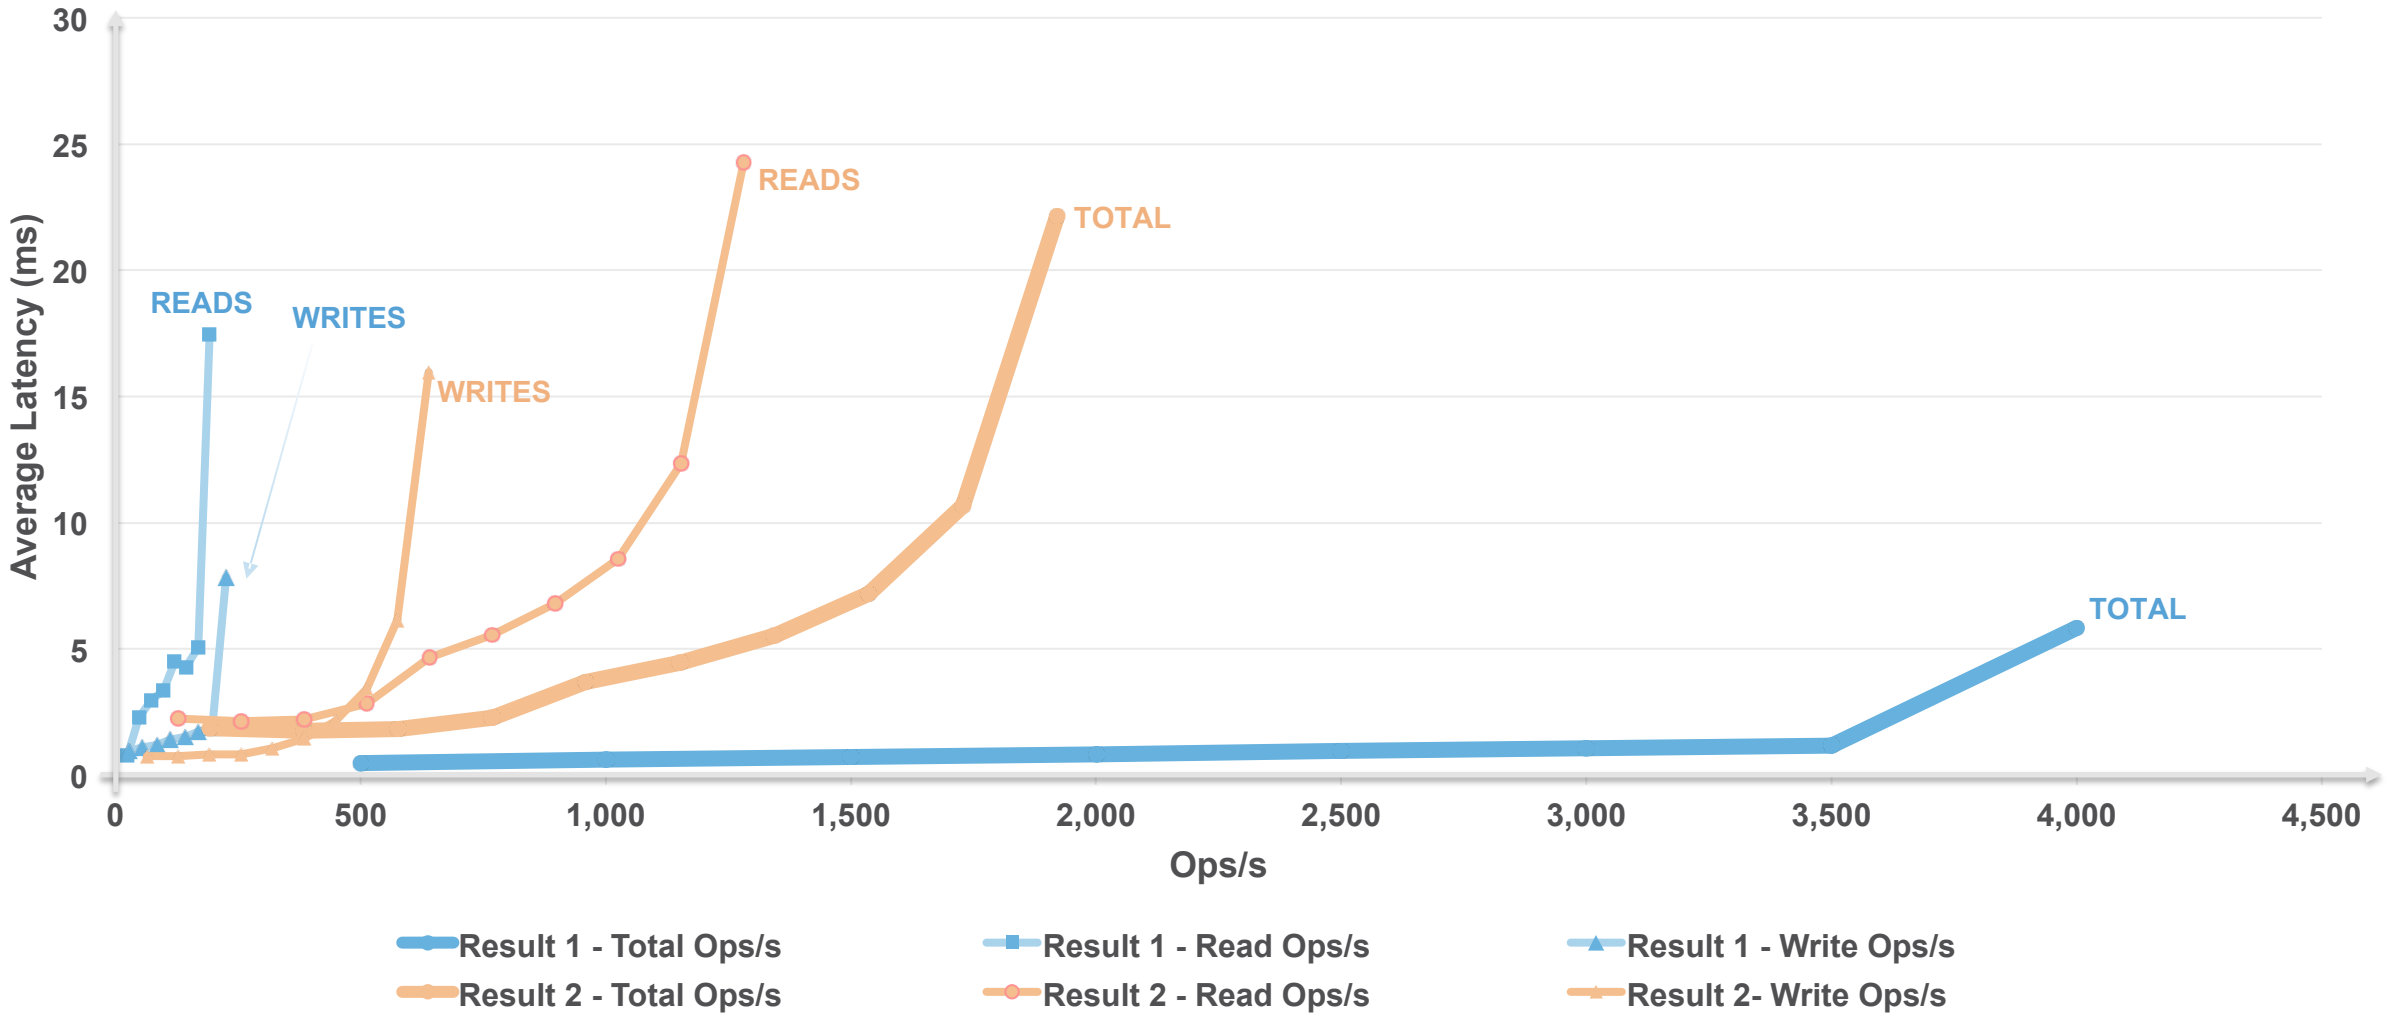

Which Is Better?

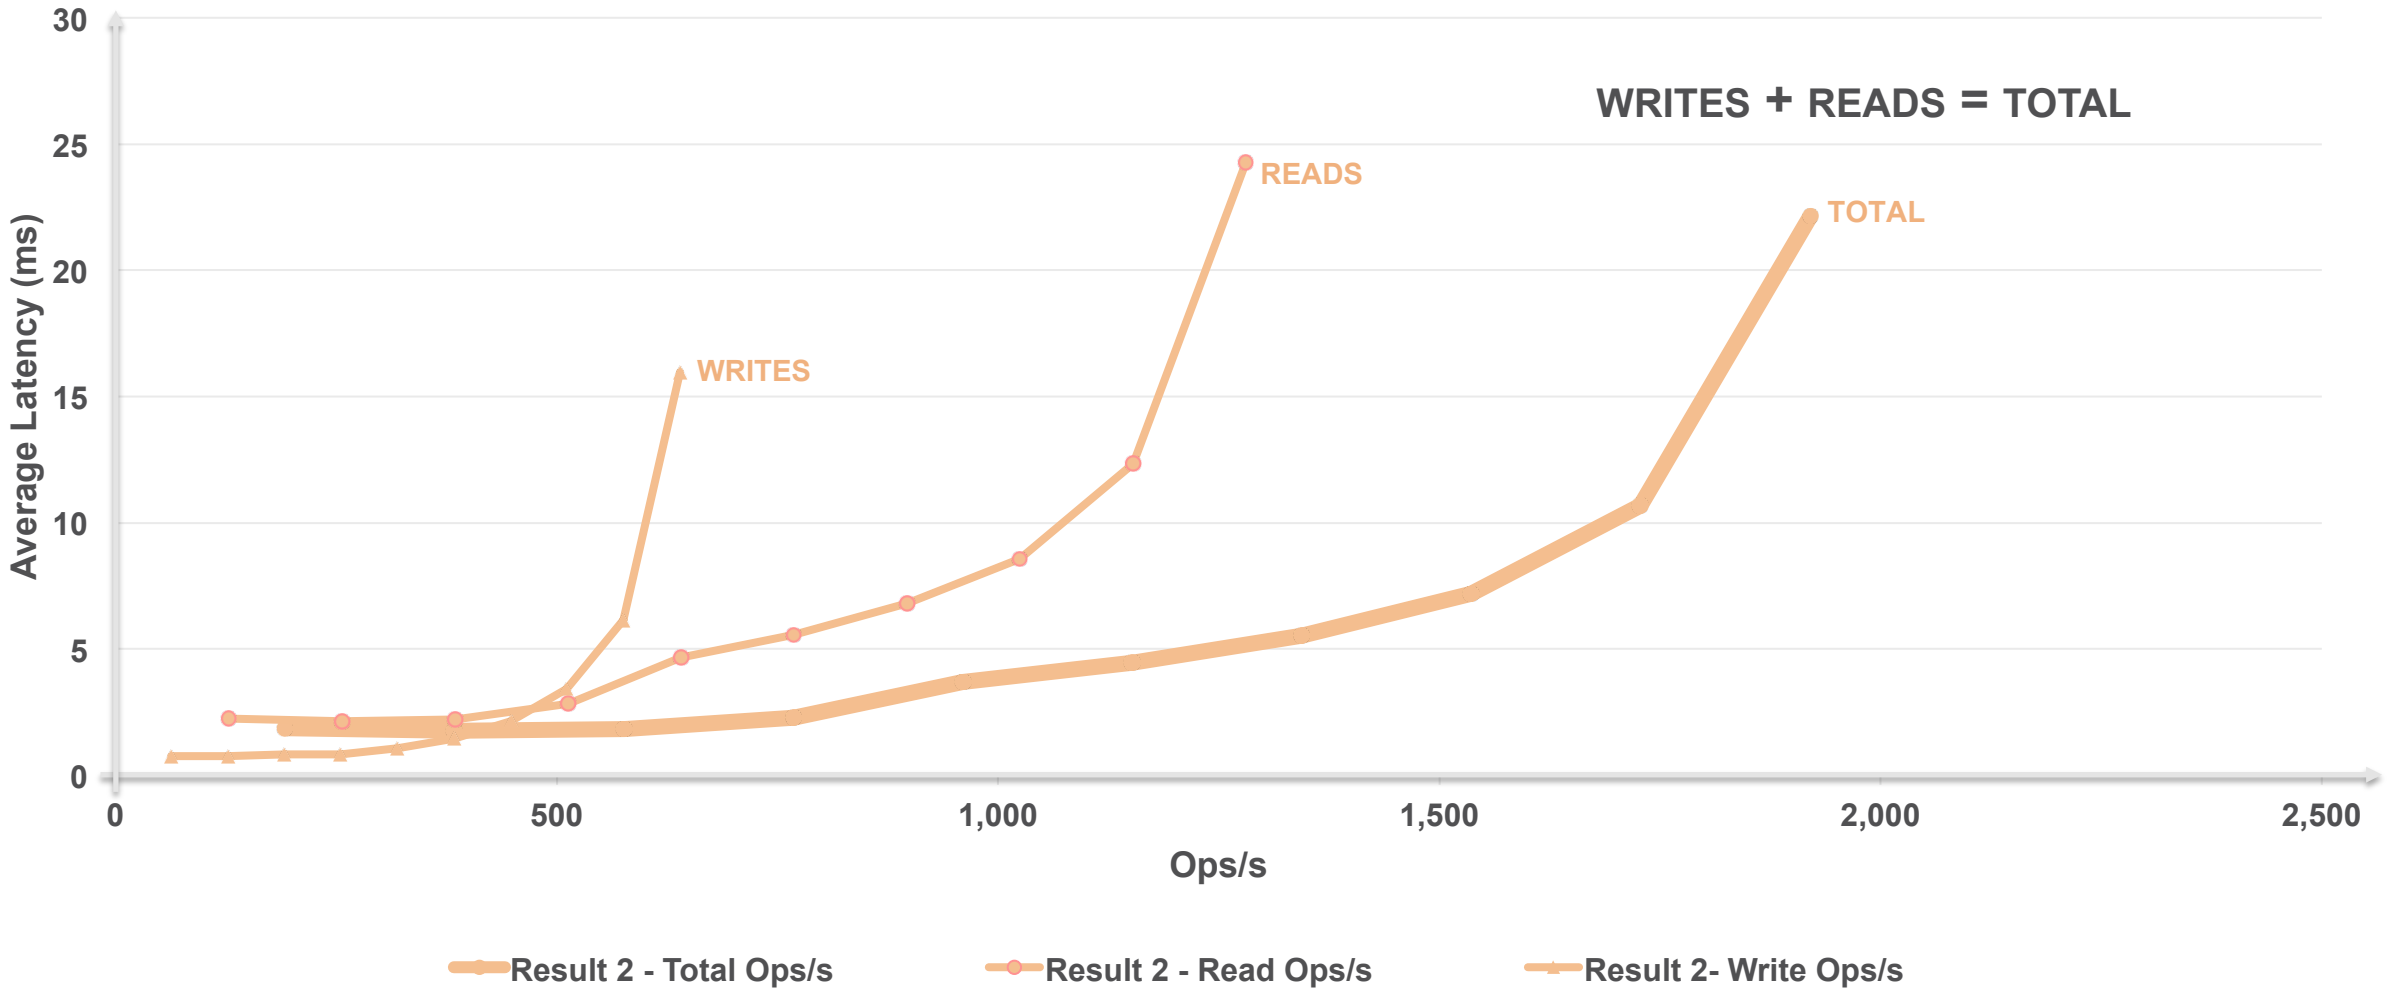

Which Is Better?

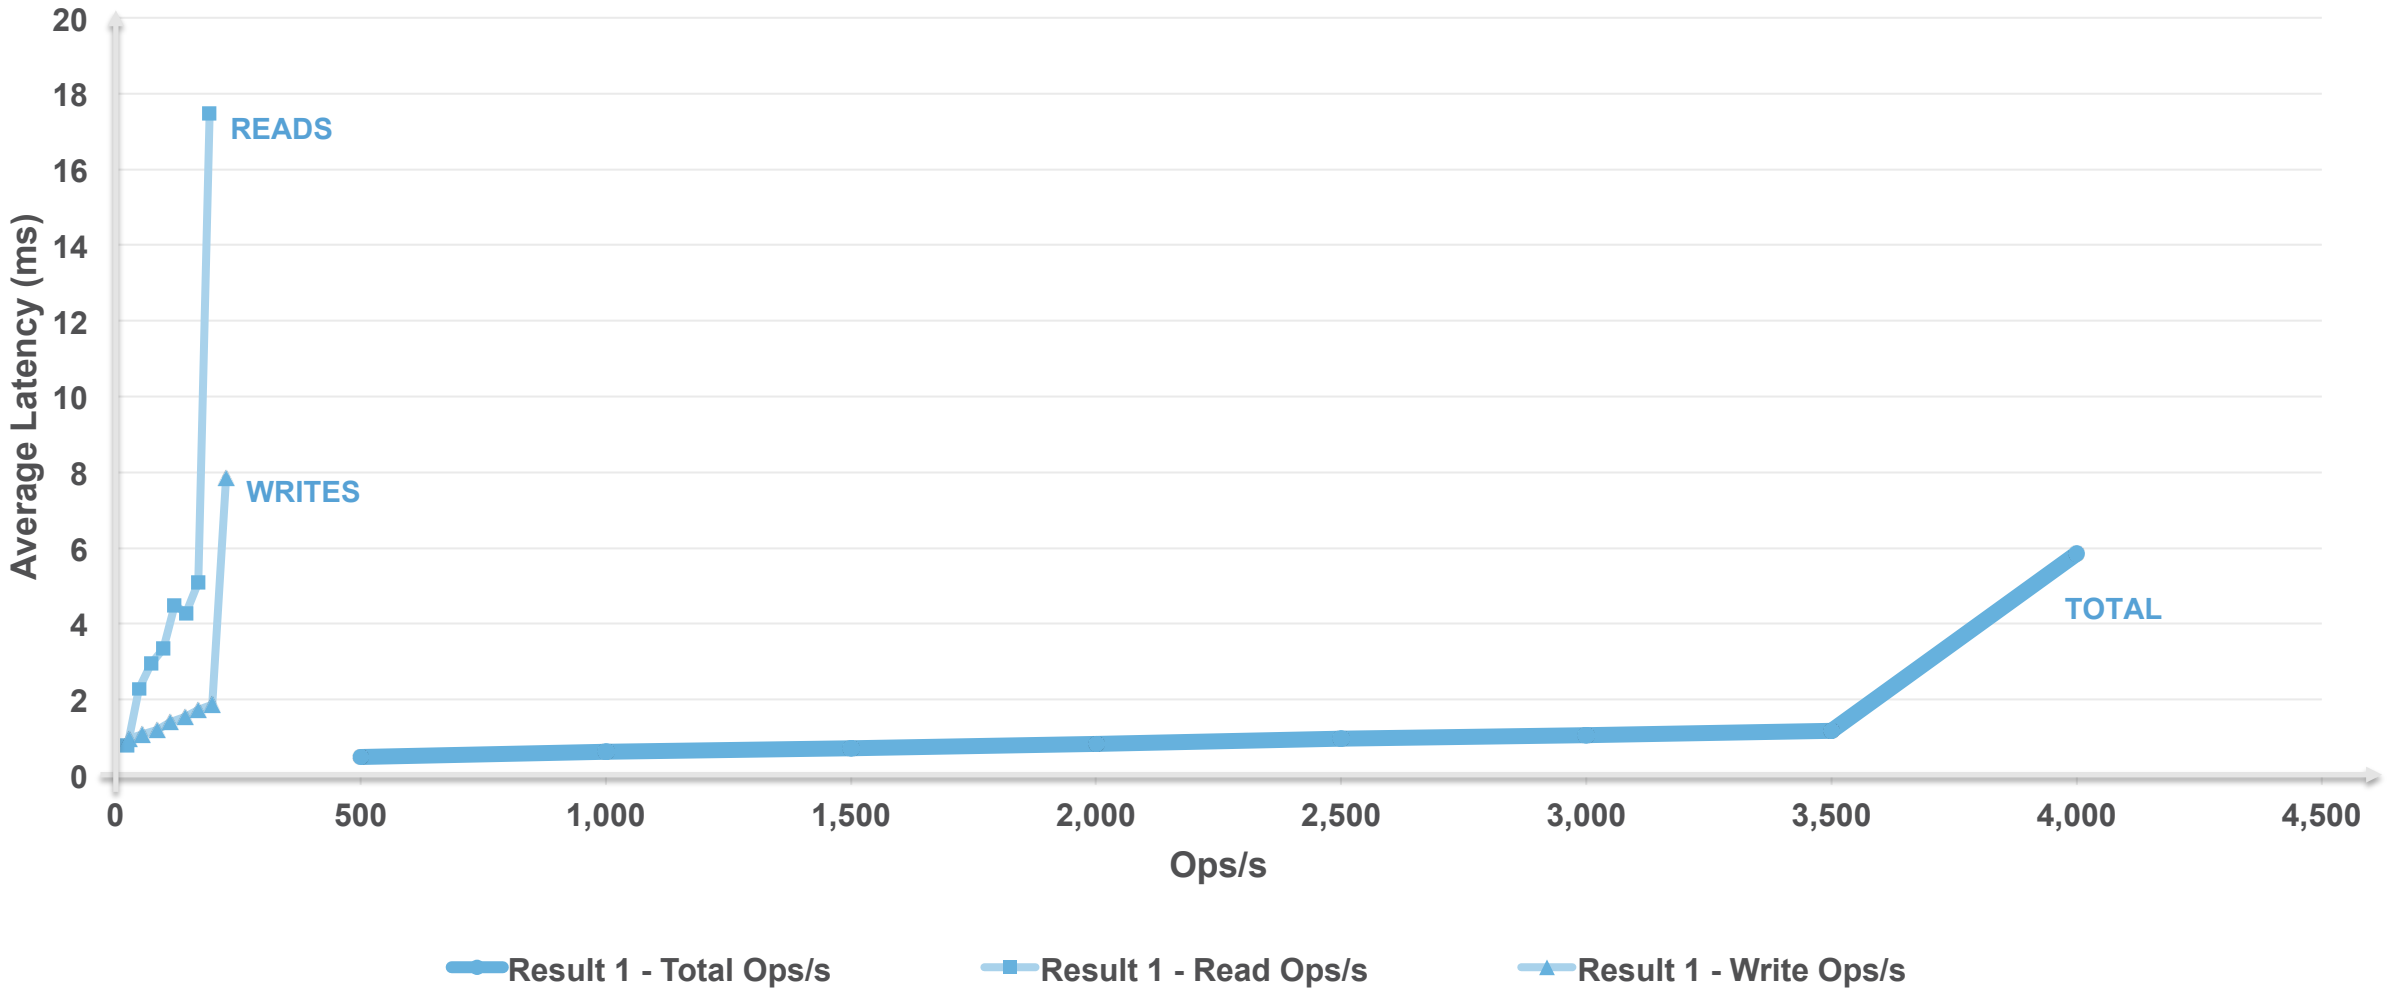

Which Is Better?

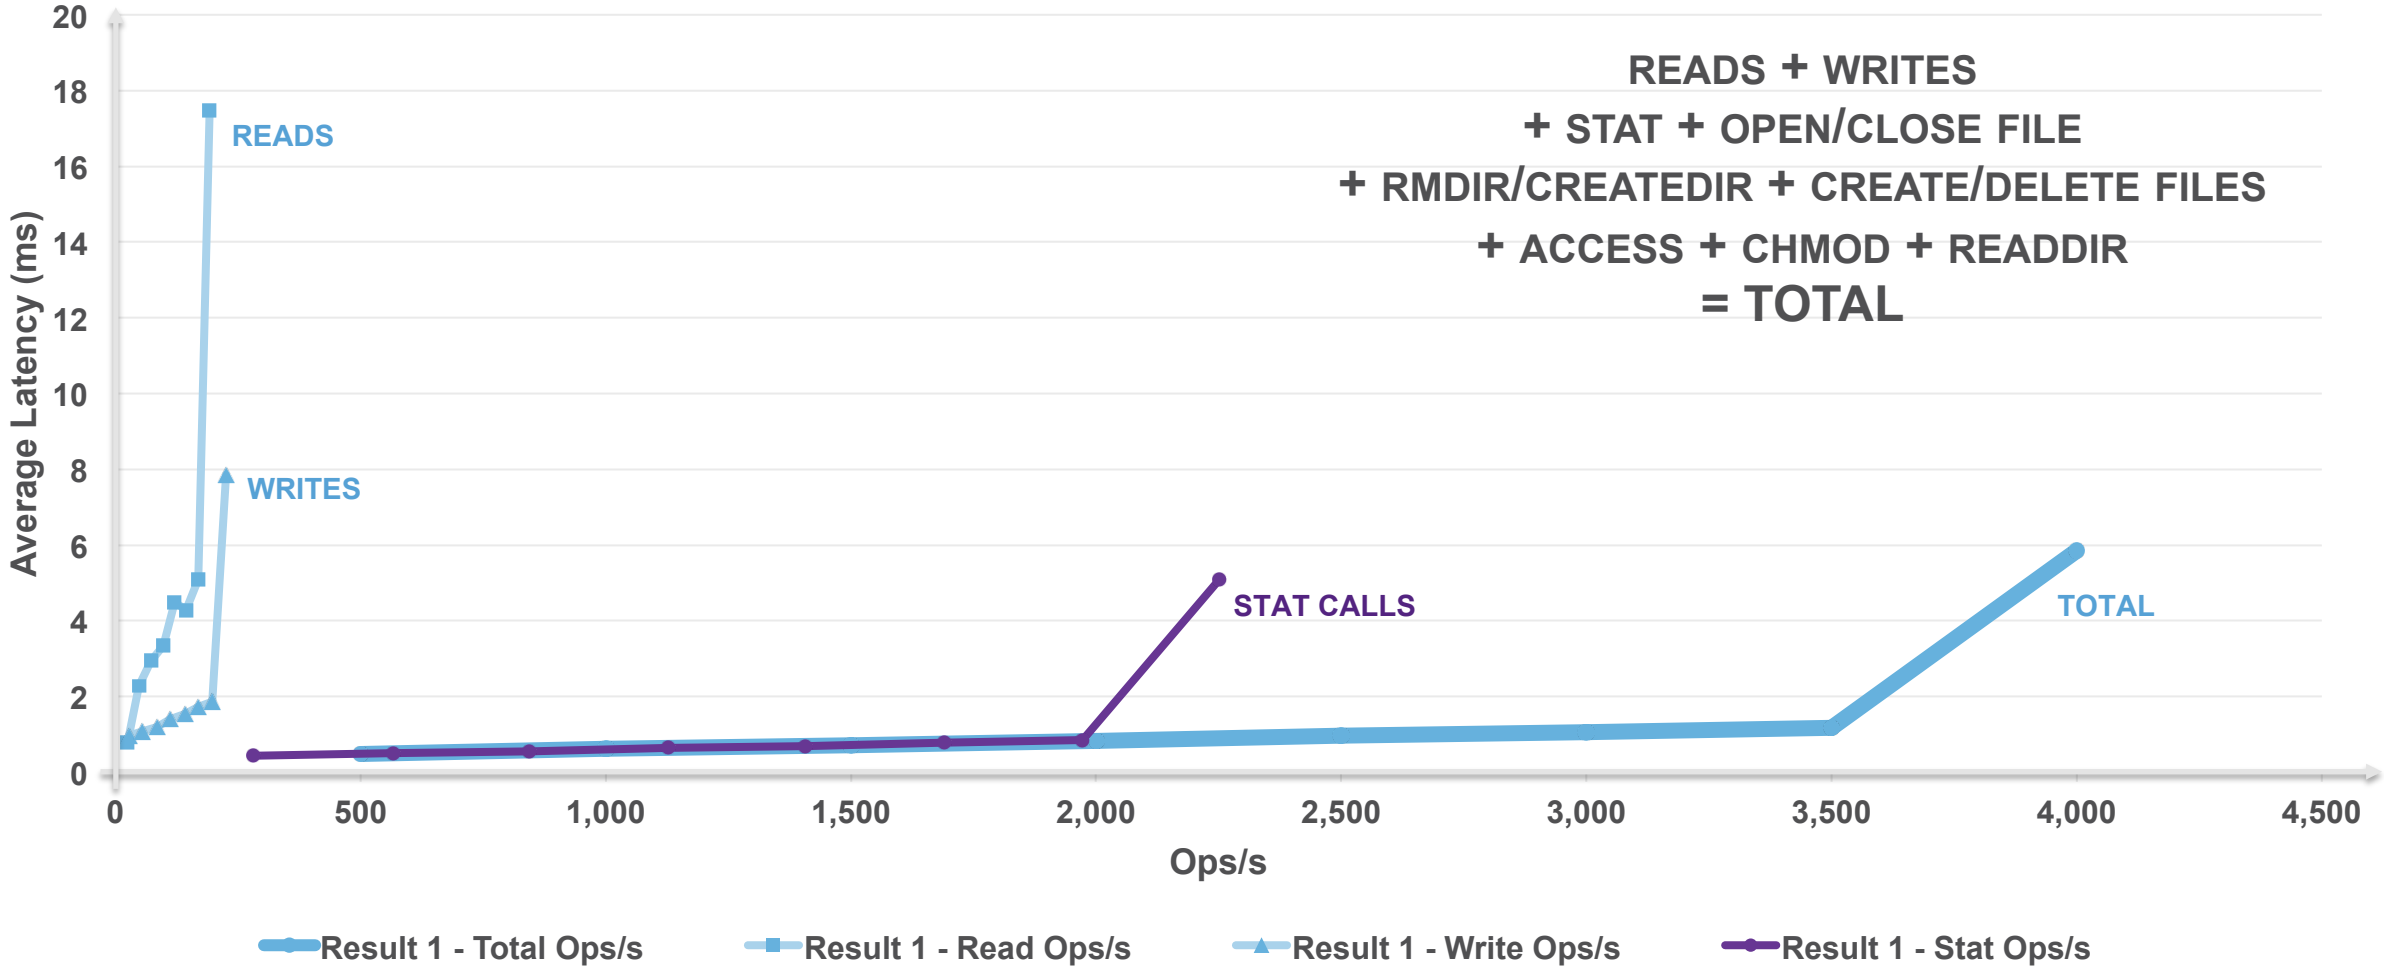

Which Is Better?

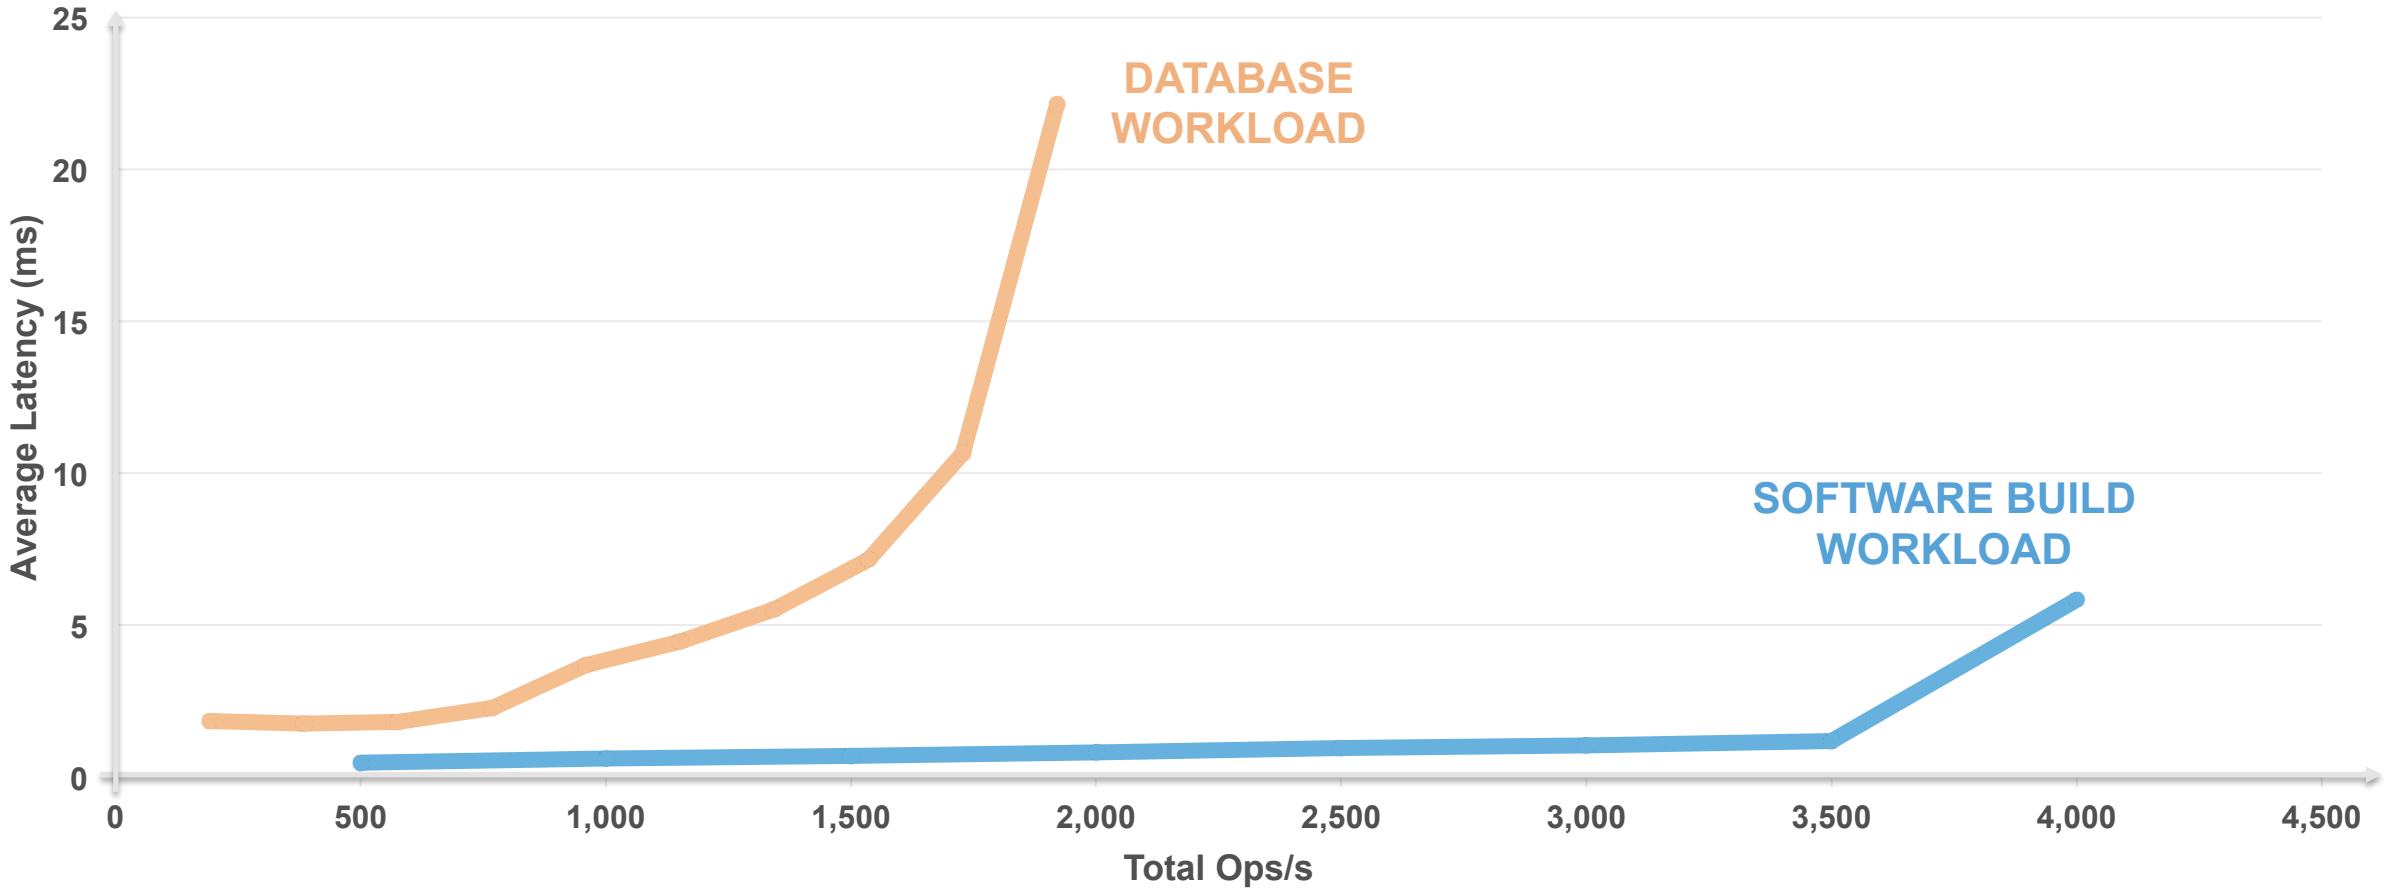

Which Is Better?

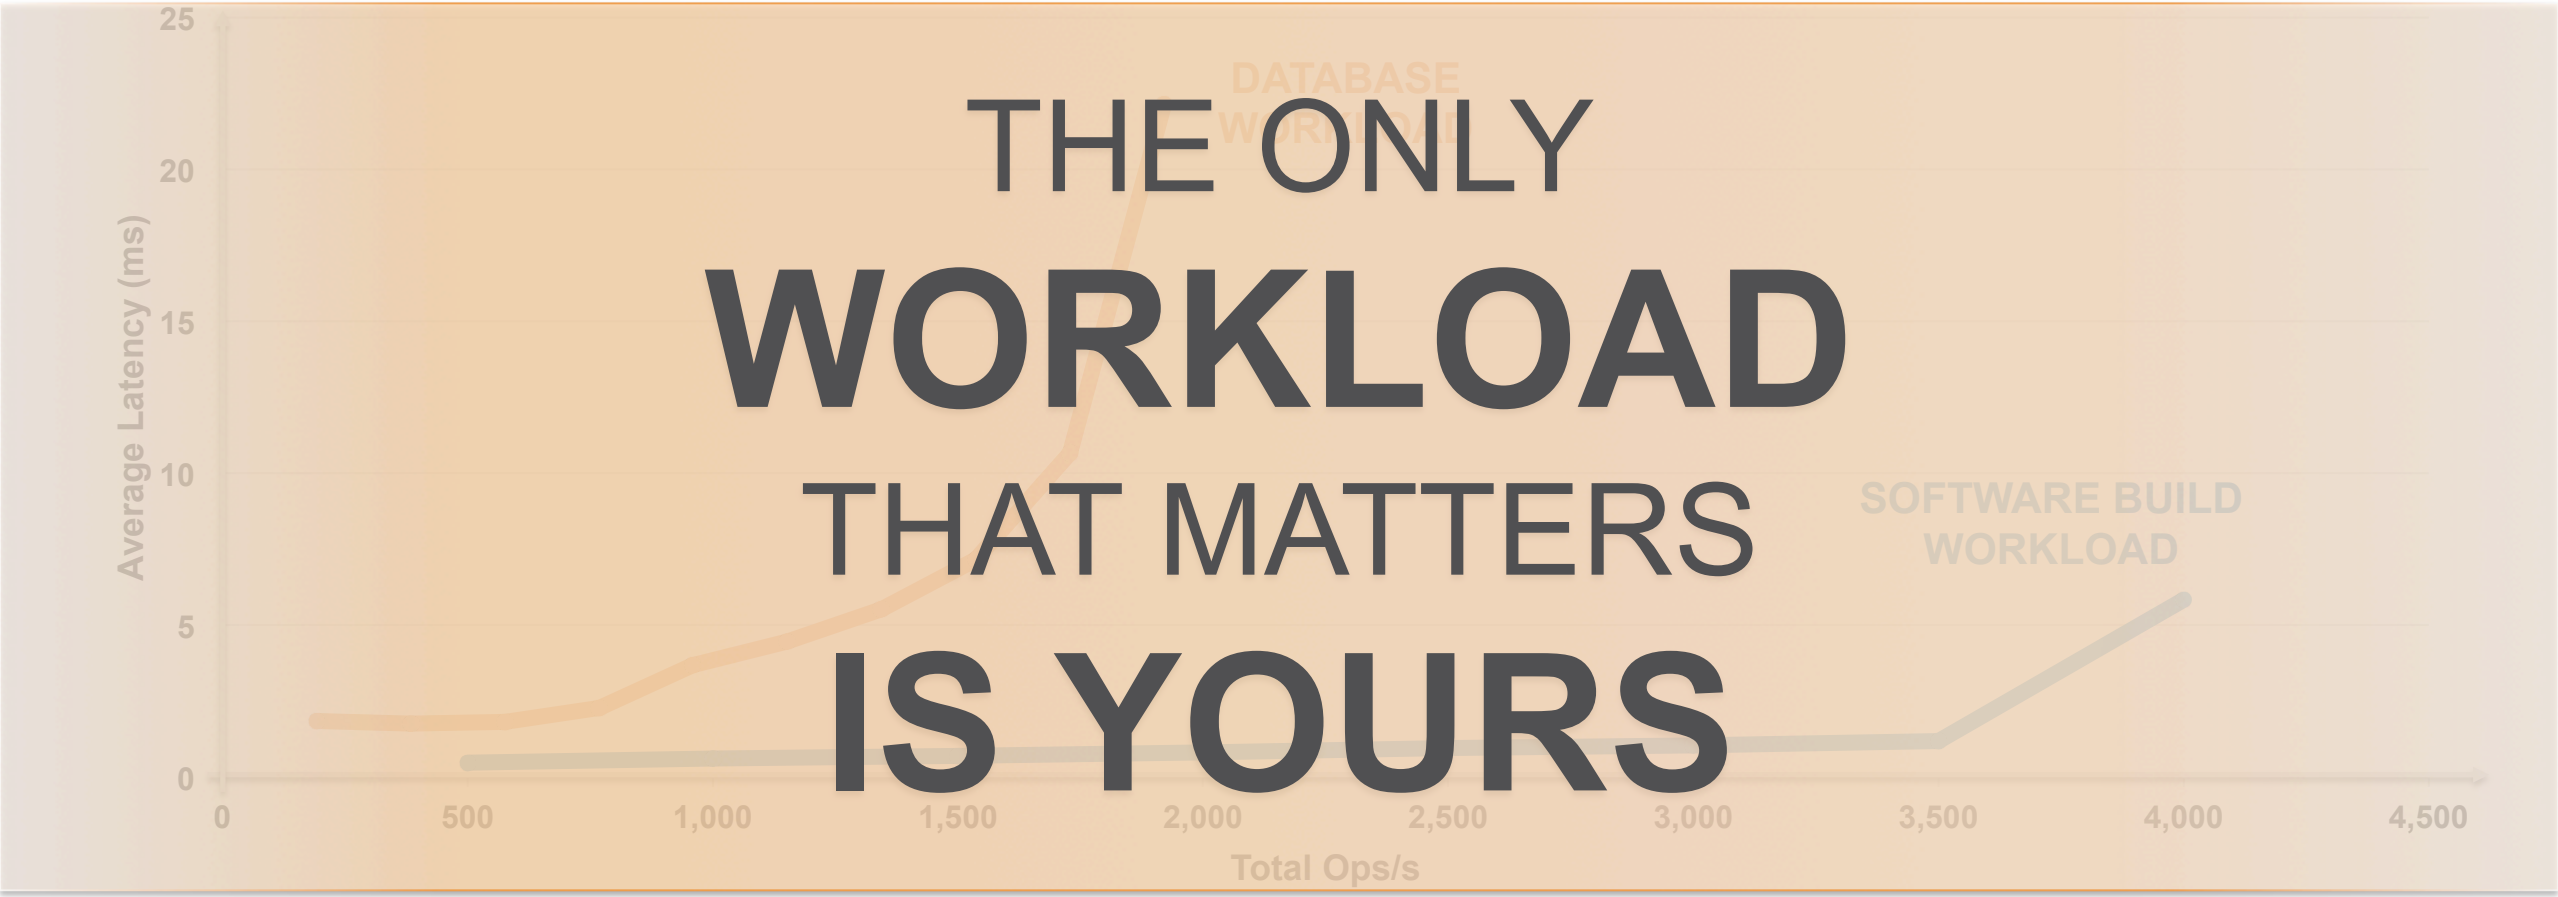

THE ONLY

WORKLOAD

THAT MATTERS

IS YOURS

Agenda

INTRO	INTRODUCTION
FILESYSTEM	FILESYSTEM LAYOUT FUNDAMENTALS
FS OPS	WHAT ARE FILESYSTEM OPERATIONS?
FILE PERF	PERFORMANCE IN A FILE ORIENTED WORLD
FUN	GRAPH FUN WITH PERFORMANCE
END	SUMMARY

Storage Performance Benchmarking

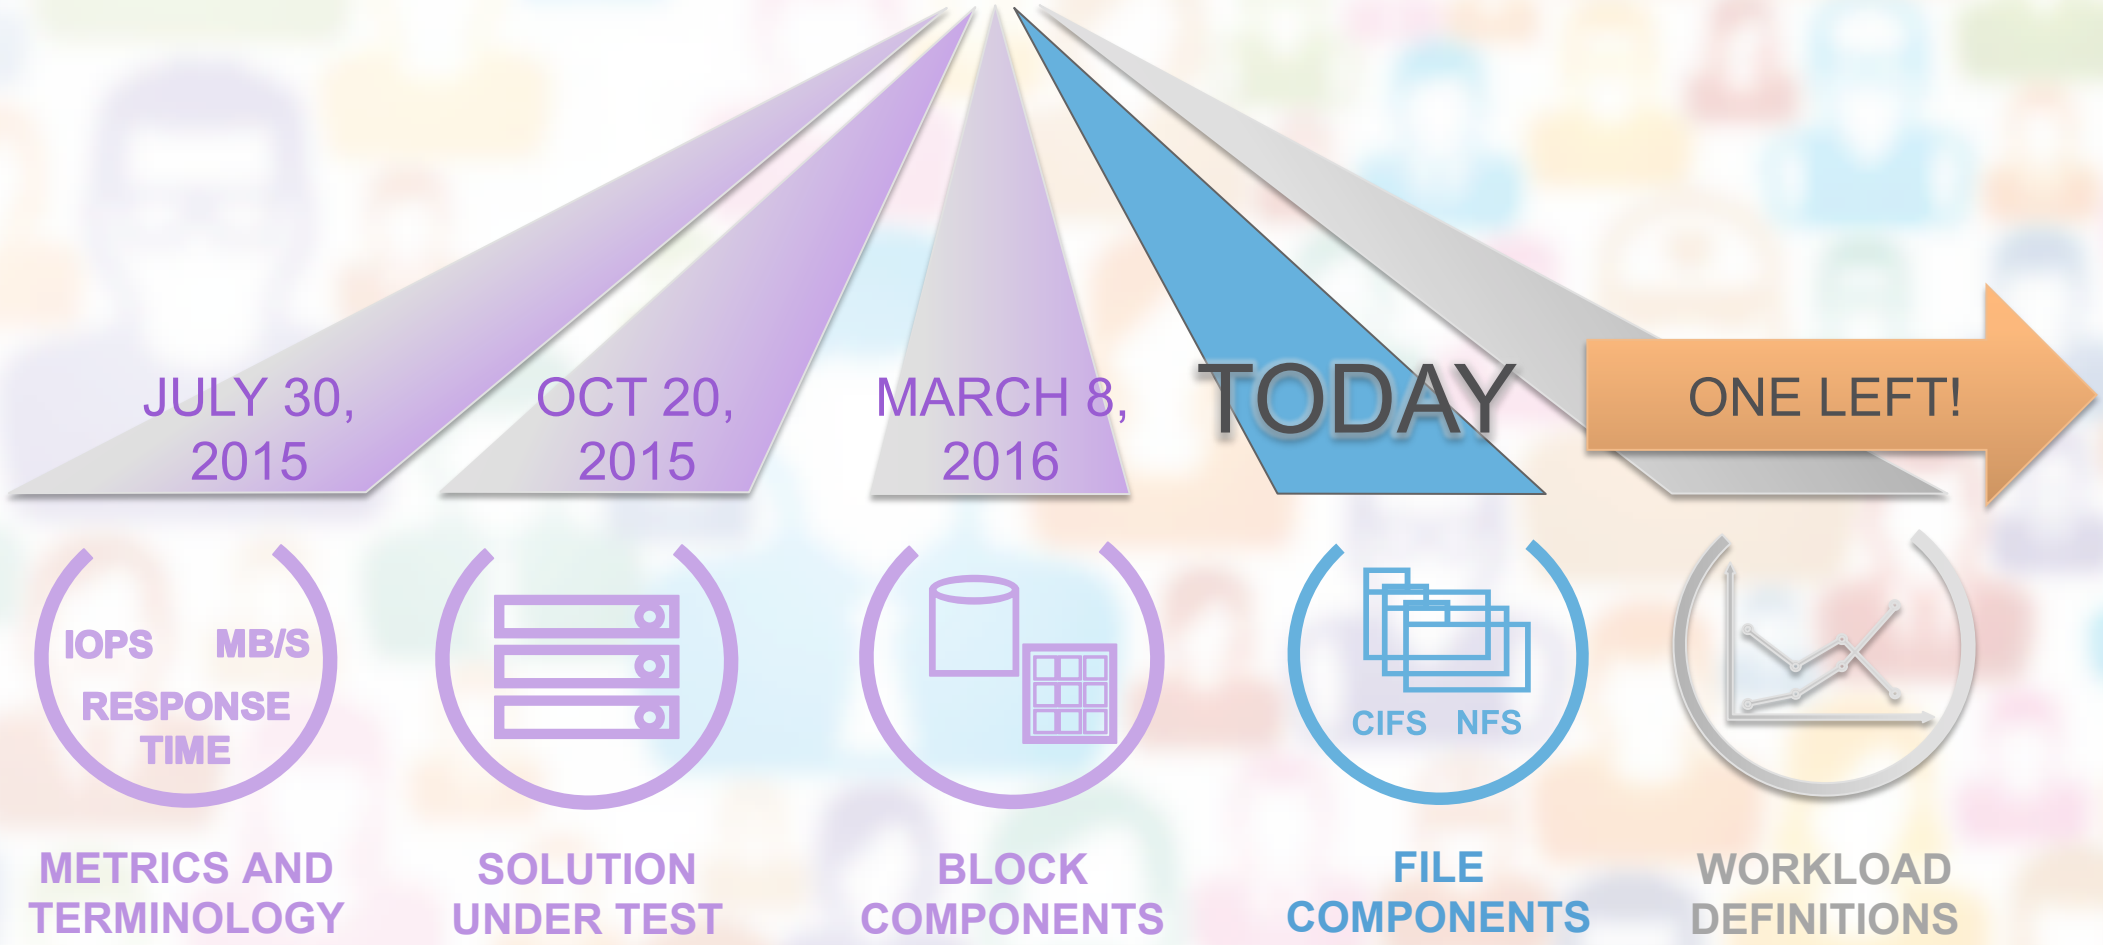

After This Webcast

SNIA ESF | ETHERNET STORAGE

- A PDF and a PPT of the slides for this and all previous parts of this Webcast series will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand
  - PPT and PDF: http://www.snia.org/forums/esf/knowledge/webcasts
  - Presentation Recording: https://www.brighttalk.com/webcast/663/219127
  - A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog
  - http://sniaesfblog.org/
- Follow us on Twitter @SNIAESF, @RogovMark, @KenCantrellJr, @alextangent, @DrJMetz
- Final Webcast – First Half of 2017
  - "Storage Performance Benchmarking: Part 5"

THANK YOU!

Client Caching Affects Operation Mixes

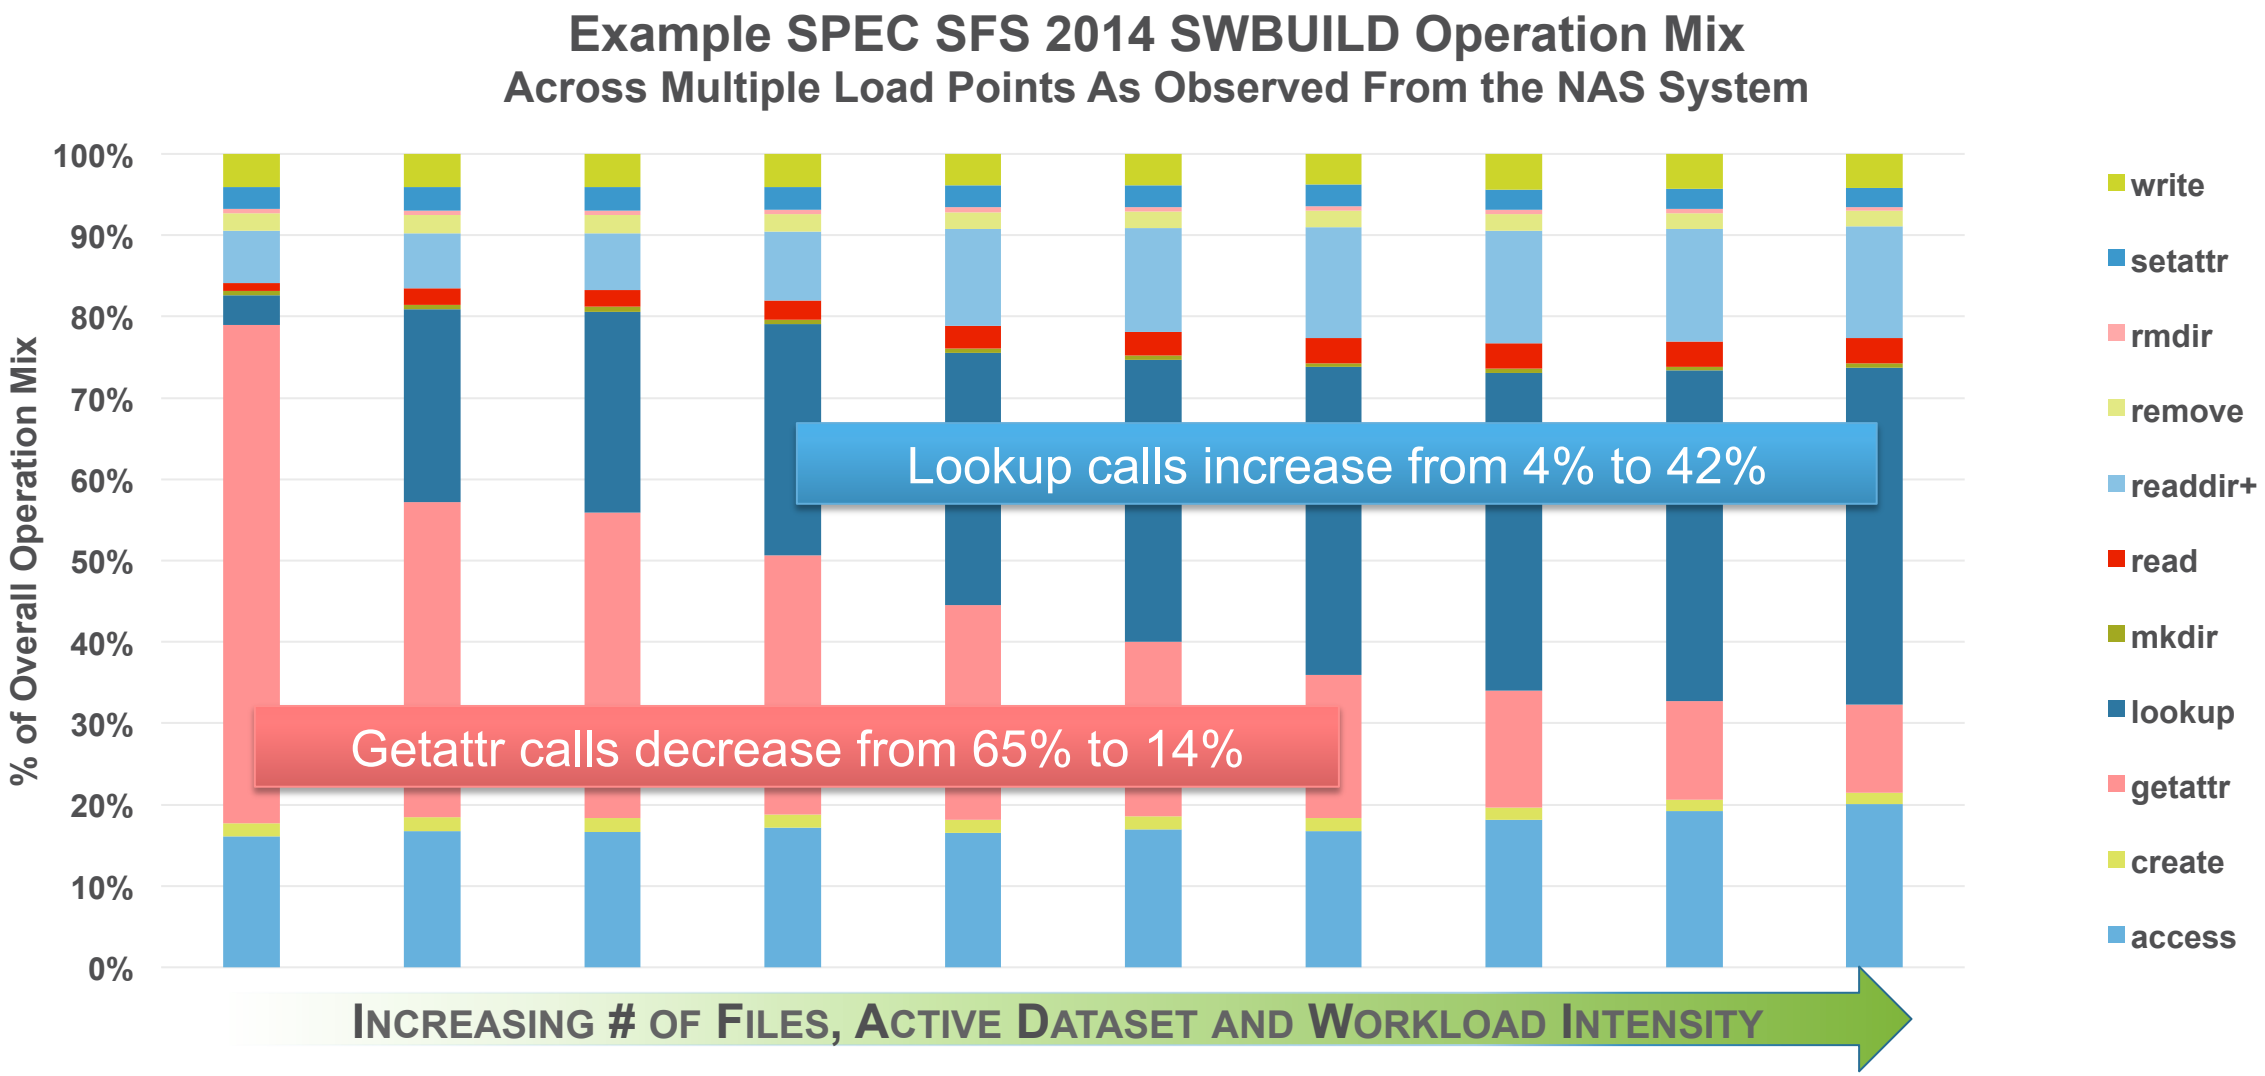

Operation Mixes For Different Applications