

#WWDC19

# Optimizing Storage in Your App

## Better performance and efficiency

Kai Kaahaina, CoreOS

Alejandro Lucena, Developer Tools

# Optimizing Storage in Your App

Better performance and efficiency

# Optimizing Storage in Your App

Better performance and efficiency

Battery life

# Optimizing Storage in Your App

Better performance and efficiency

Battery life

Performance

# Optimizing Storage in Your App

Better performance and efficiency

Battery life

Performance

Reduce footprint

# Optimizing Storage in Your App

Better performance and efficiency

Battery life

Performance

Reduce footprint

Device health

Efficient image assets

File system metadata

Syncing to disk

Serialized data files

Core Data

SQLite

# Efficient Image Assets



# Image Assets

Growing in size

Increasing screen resolutions

- iPhone 6s — 1334 x 750
- iPhone Xs — 2436 x 1125

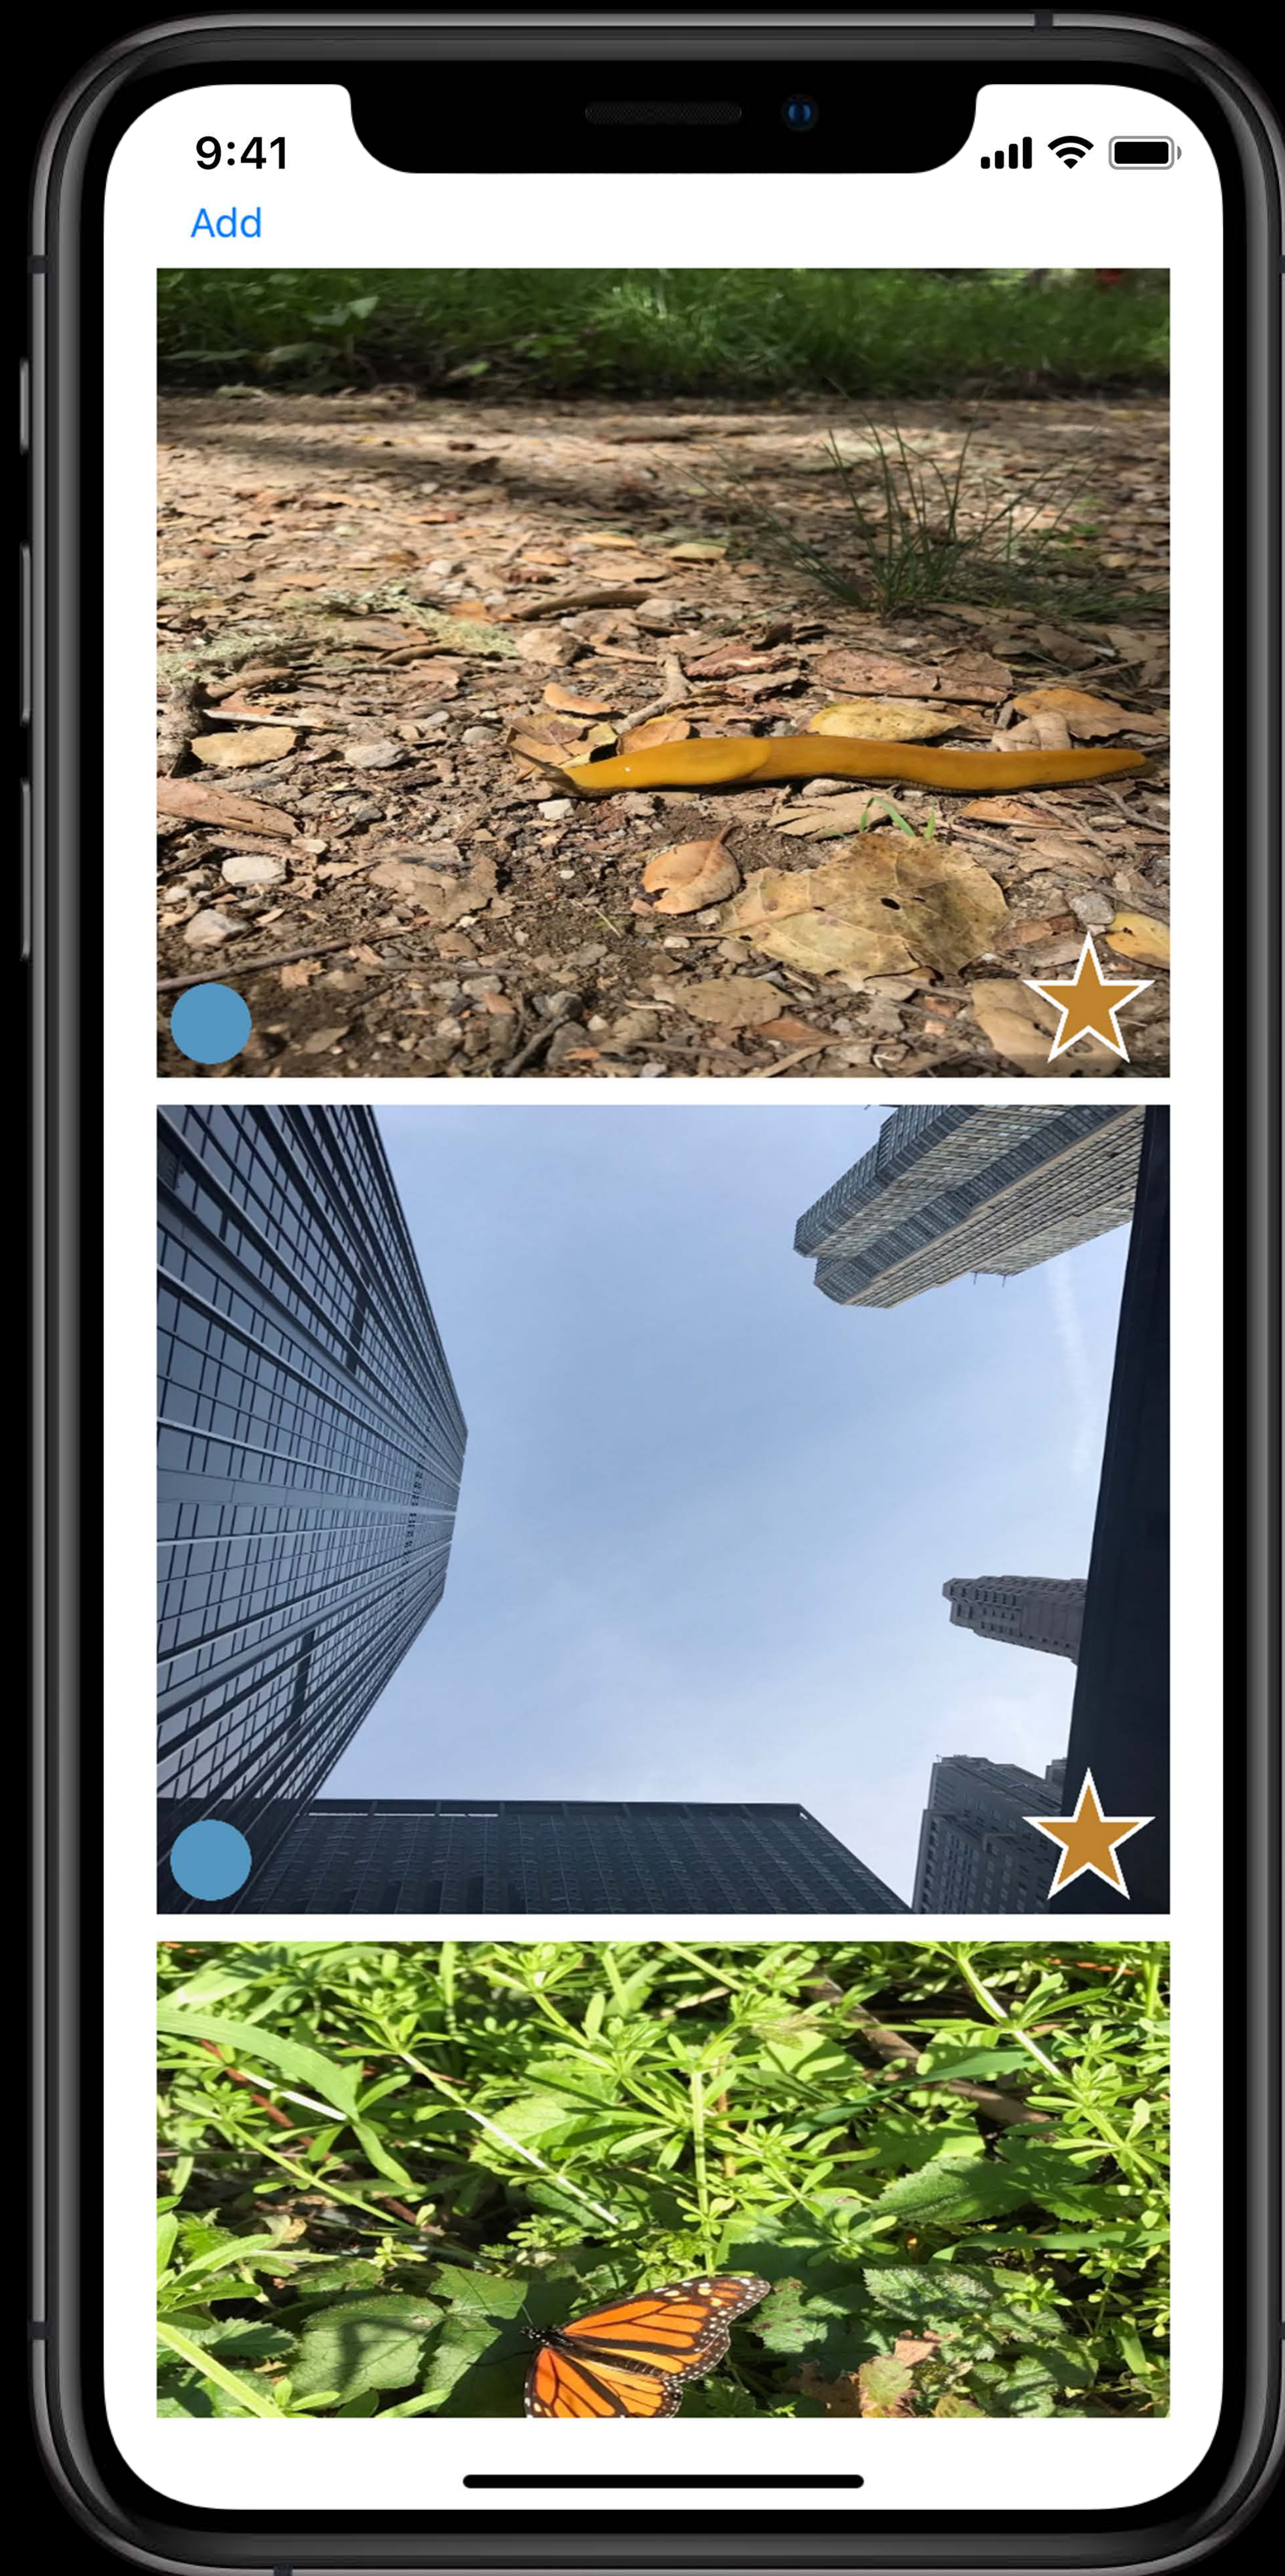
Higher resolution cameras

Harder to manage



# Image Assets

Growing in size

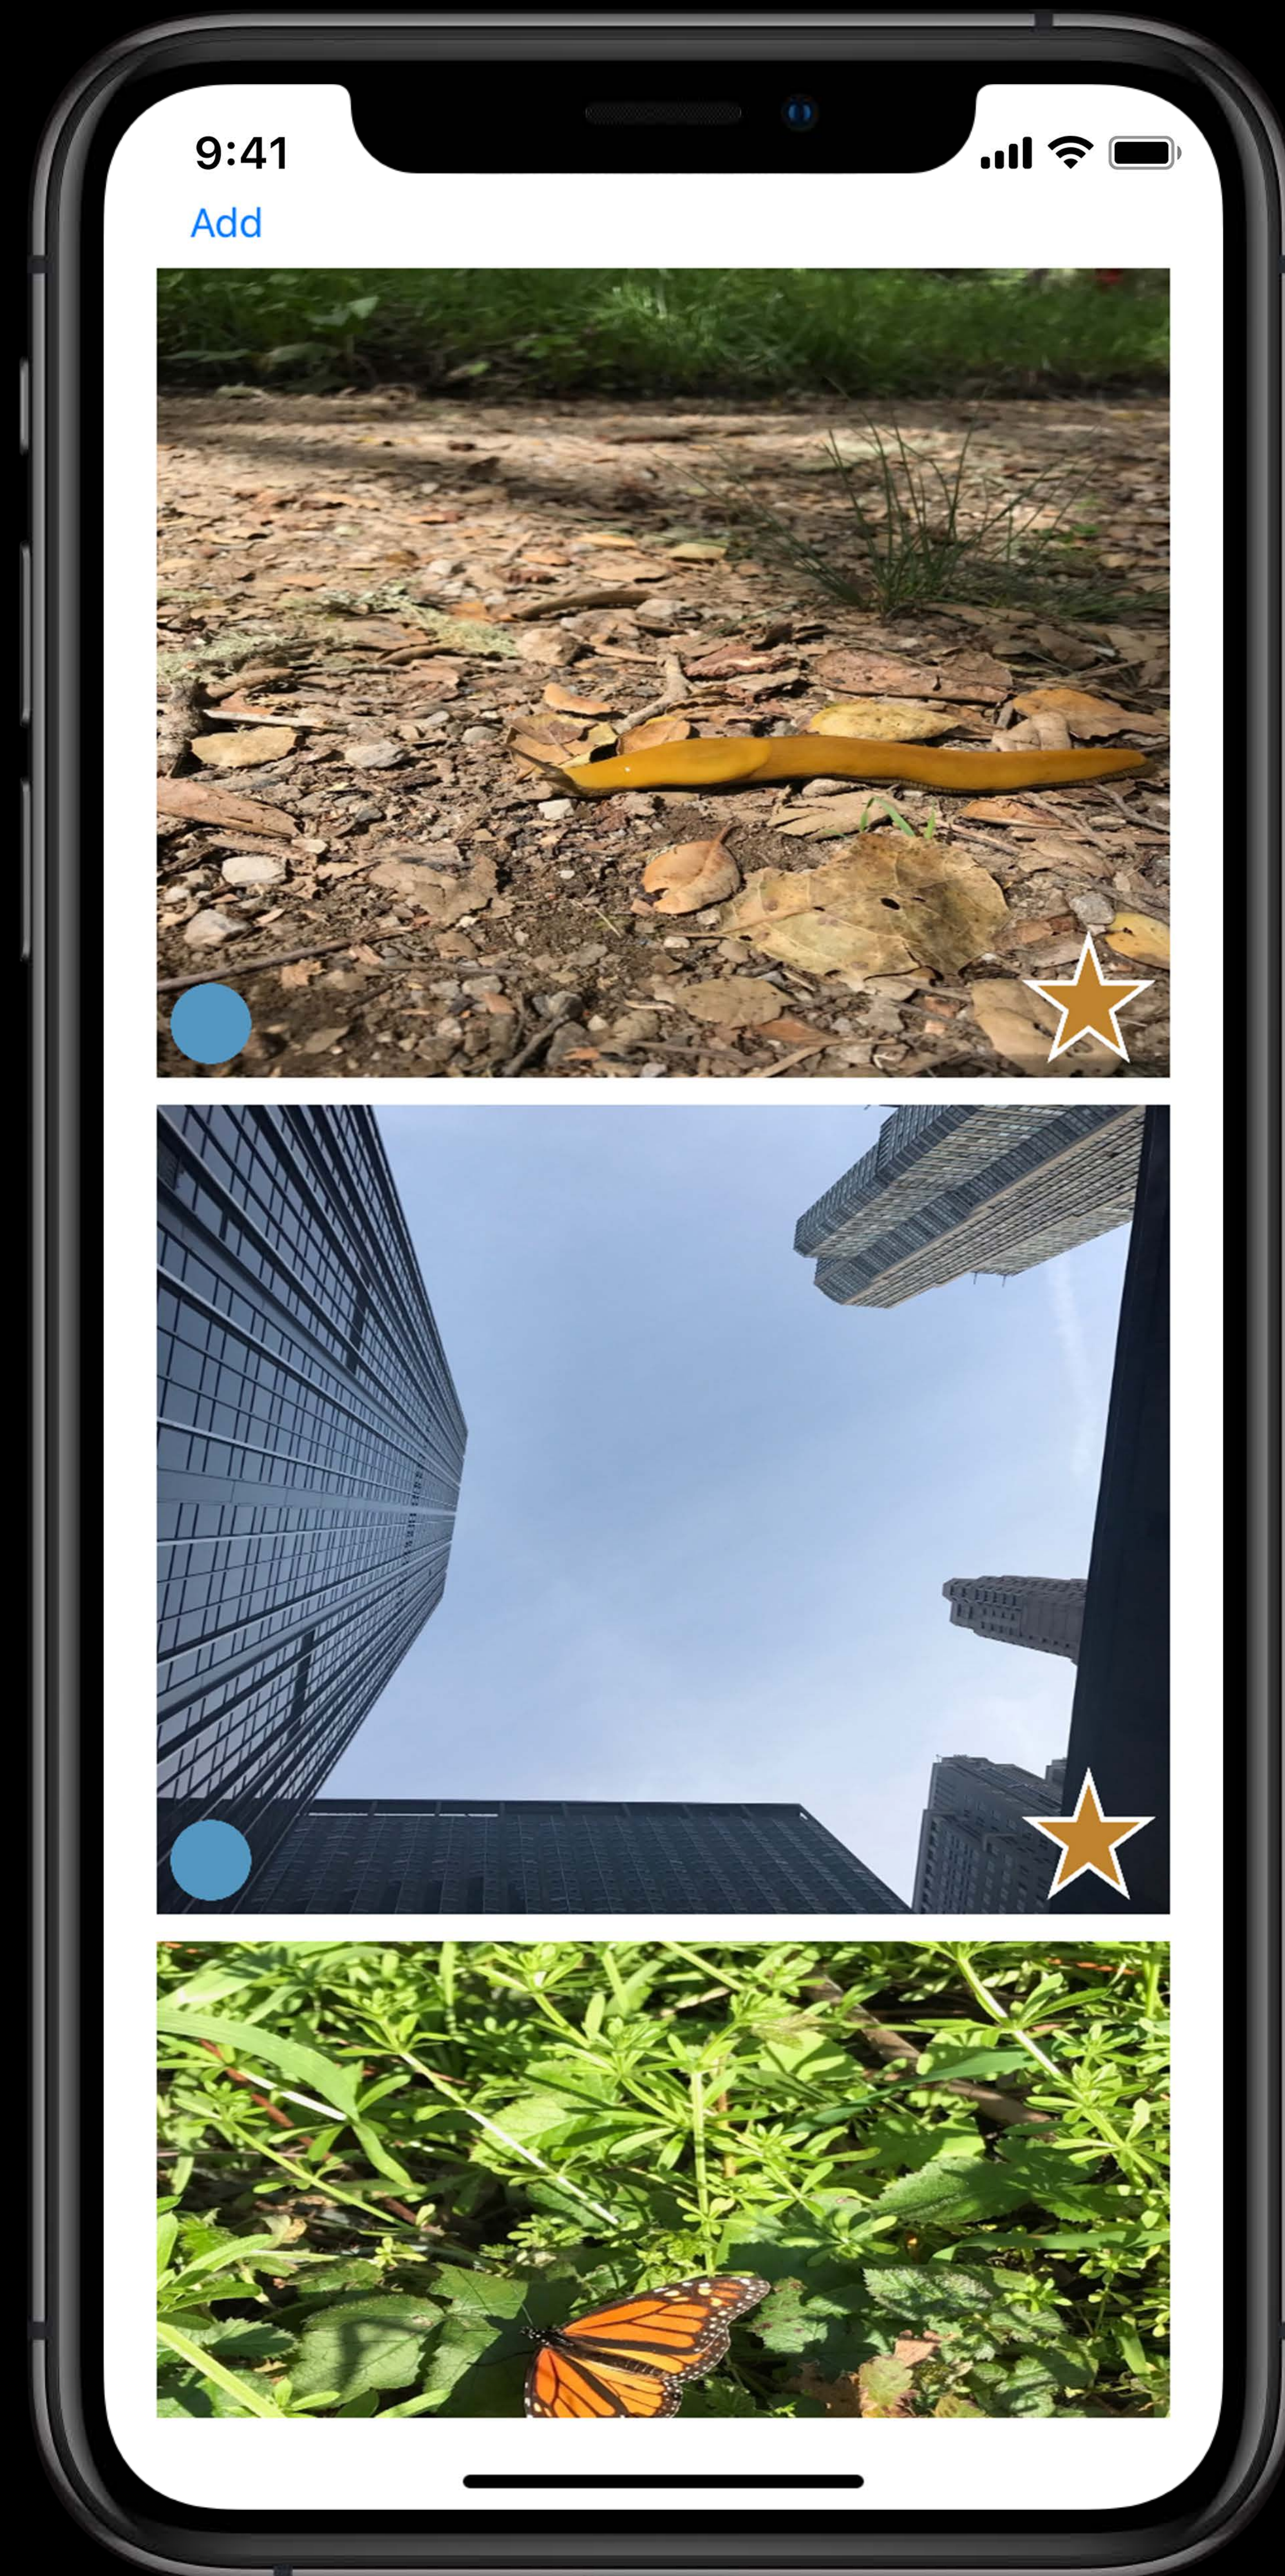




# Image Assets

Growing in size

JPEG Size: 24.6MB





# HEIC

A more efficient and capable alternative to JPEG



# HEIC

A more efficient and capable alternative to JPEG

~50% smaller files than JPEG at comparable quality





# HEIC

A more efficient and capable alternative to JPEG

~50% smaller files than JPEG at comparable quality

Smaller on disk footprint



# HEIC

A more efficient and capable alternative to JPEG

~50% smaller files than JPEG at comparable quality

Smaller on disk footprint

Faster network downloads and uploads





# HEIC

A more efficient and capable alternative to JPEG

~50% smaller files than JPEG at comparable quality

Smaller on disk footprint

Faster network downloads and uploads

Faster to load and save to disk





# HEIC

A more efficient and capable alternative to JPEG

Store auxiliary images (depth, disparity, and so on)



# HEIC

A more efficient and capable alternative to JPEG

Store auxiliary images (depth, disparity, and so on)

Supports alpha





# HEIC

A more efficient and capable alternative to JPEG

Store auxiliary images (depth, disparity, and so on)

Supports alpha

Lossless compression



# HEIC

A more efficient and capable alternative to JPEG

Store auxiliary images (depth, disparity, and so on)

Supports alpha

Lossless compression

Multiple images in a single container





# HEIC

A more efficient and capable alternative to JPEG

Available starting with

- iOS 11
- macOS High Sierra

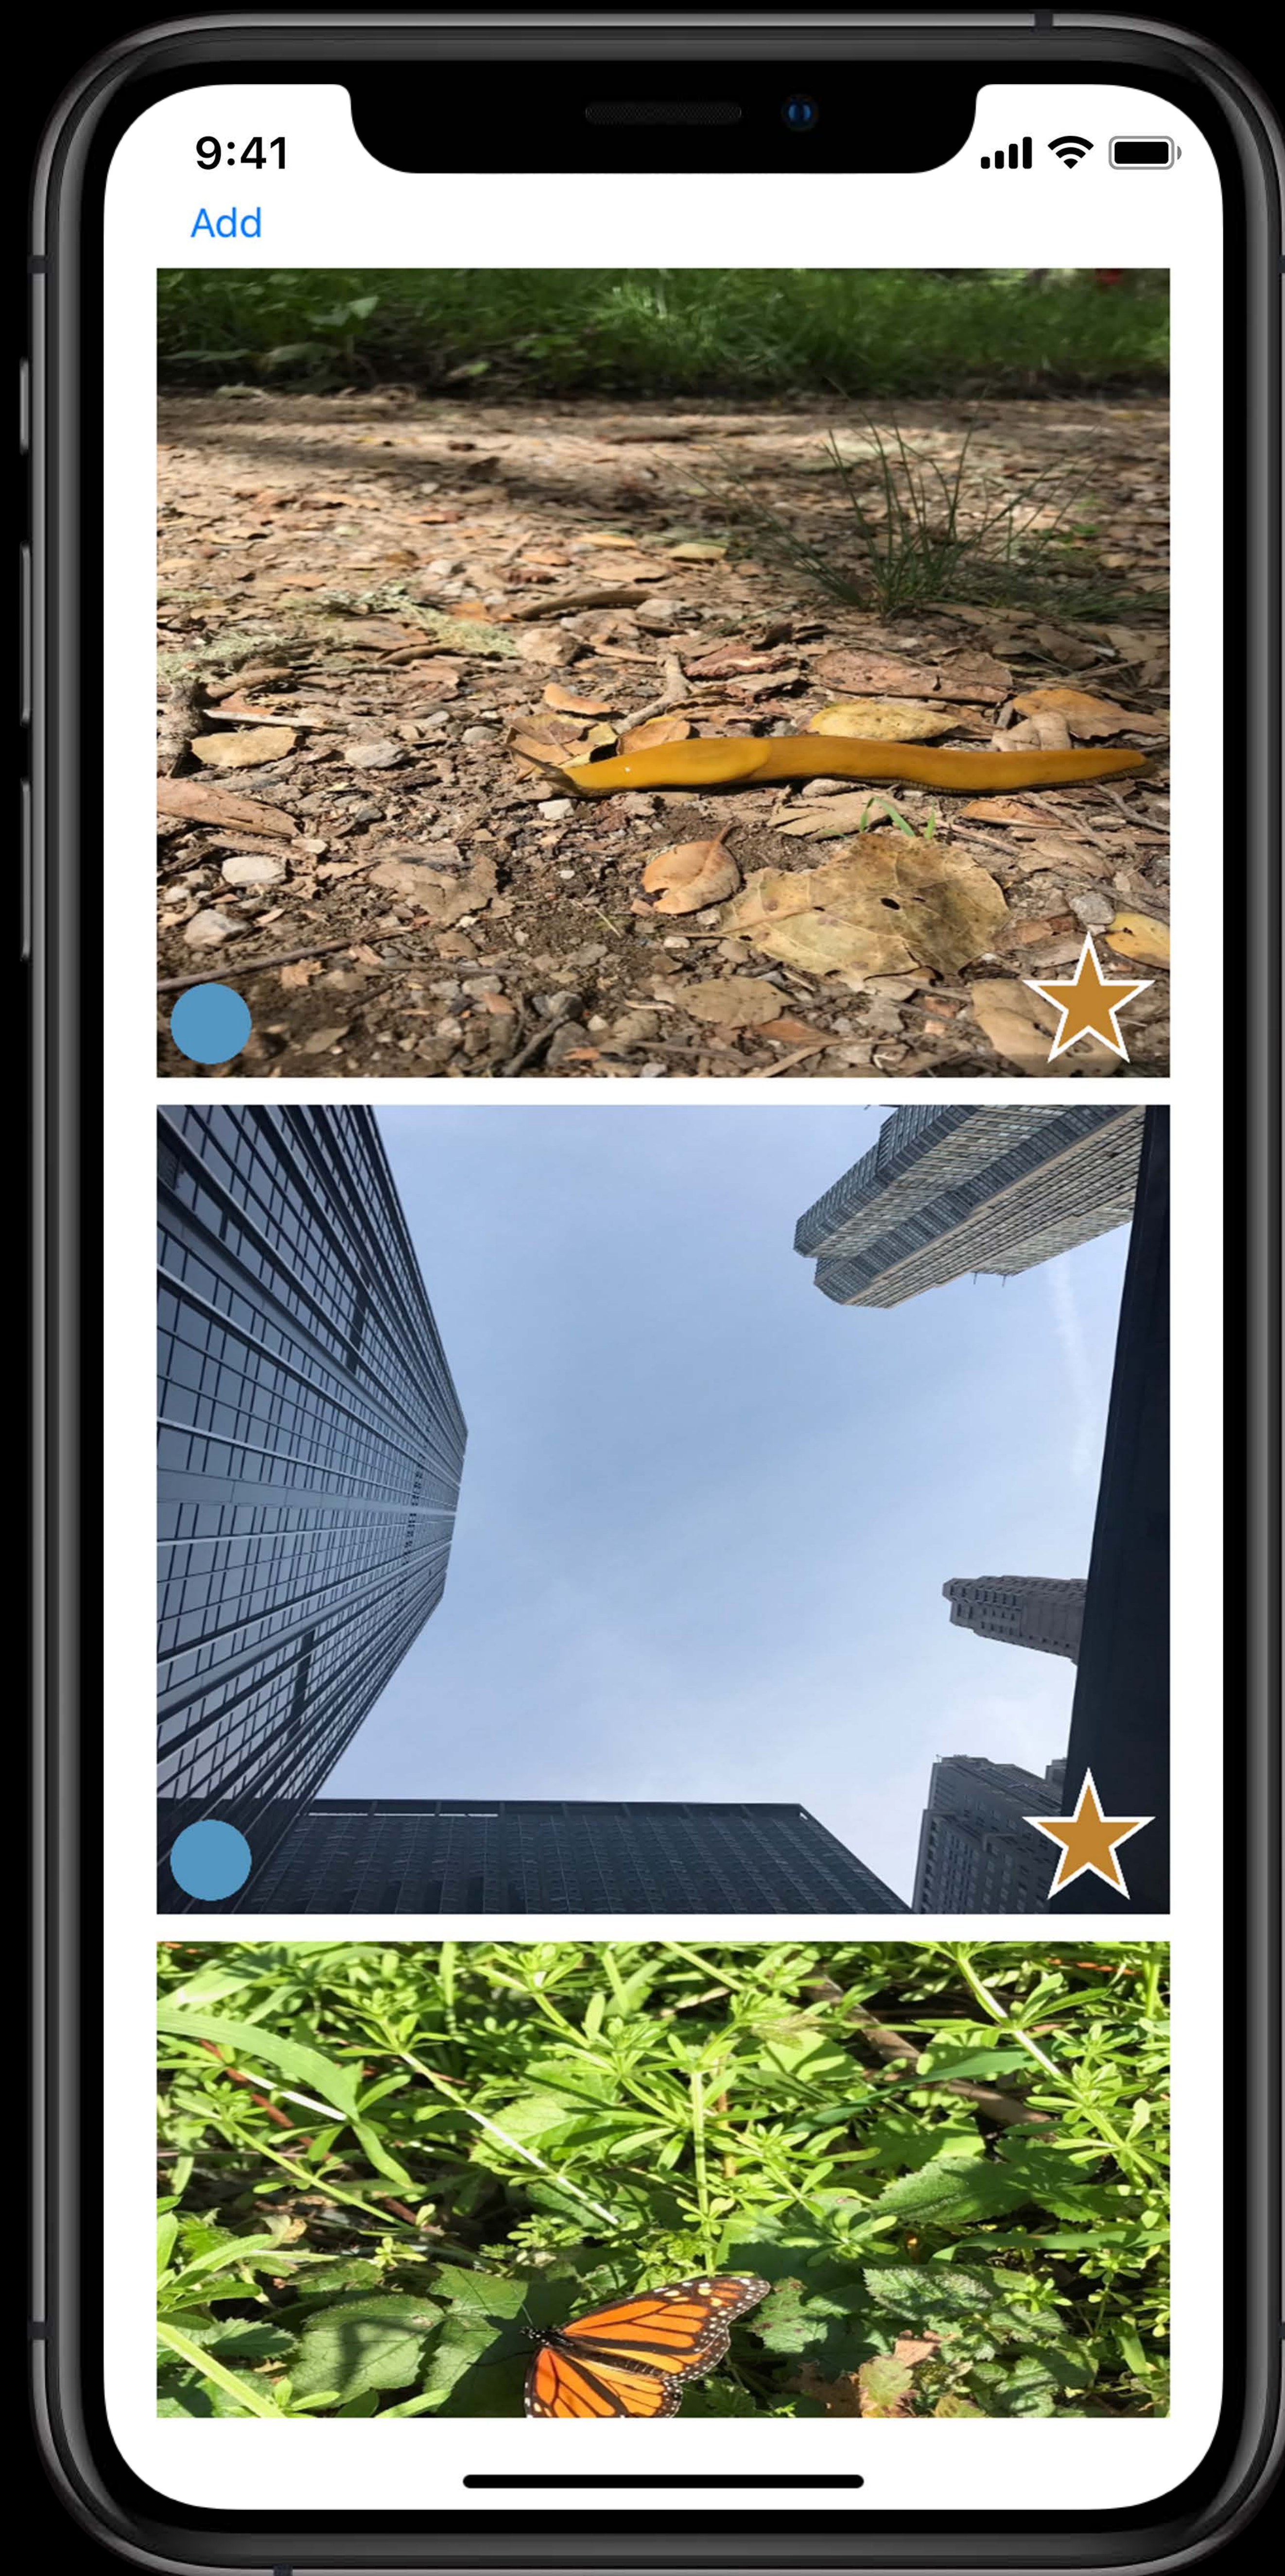




# Image Assets

Growing in size

JPEG Size: 24.6MB



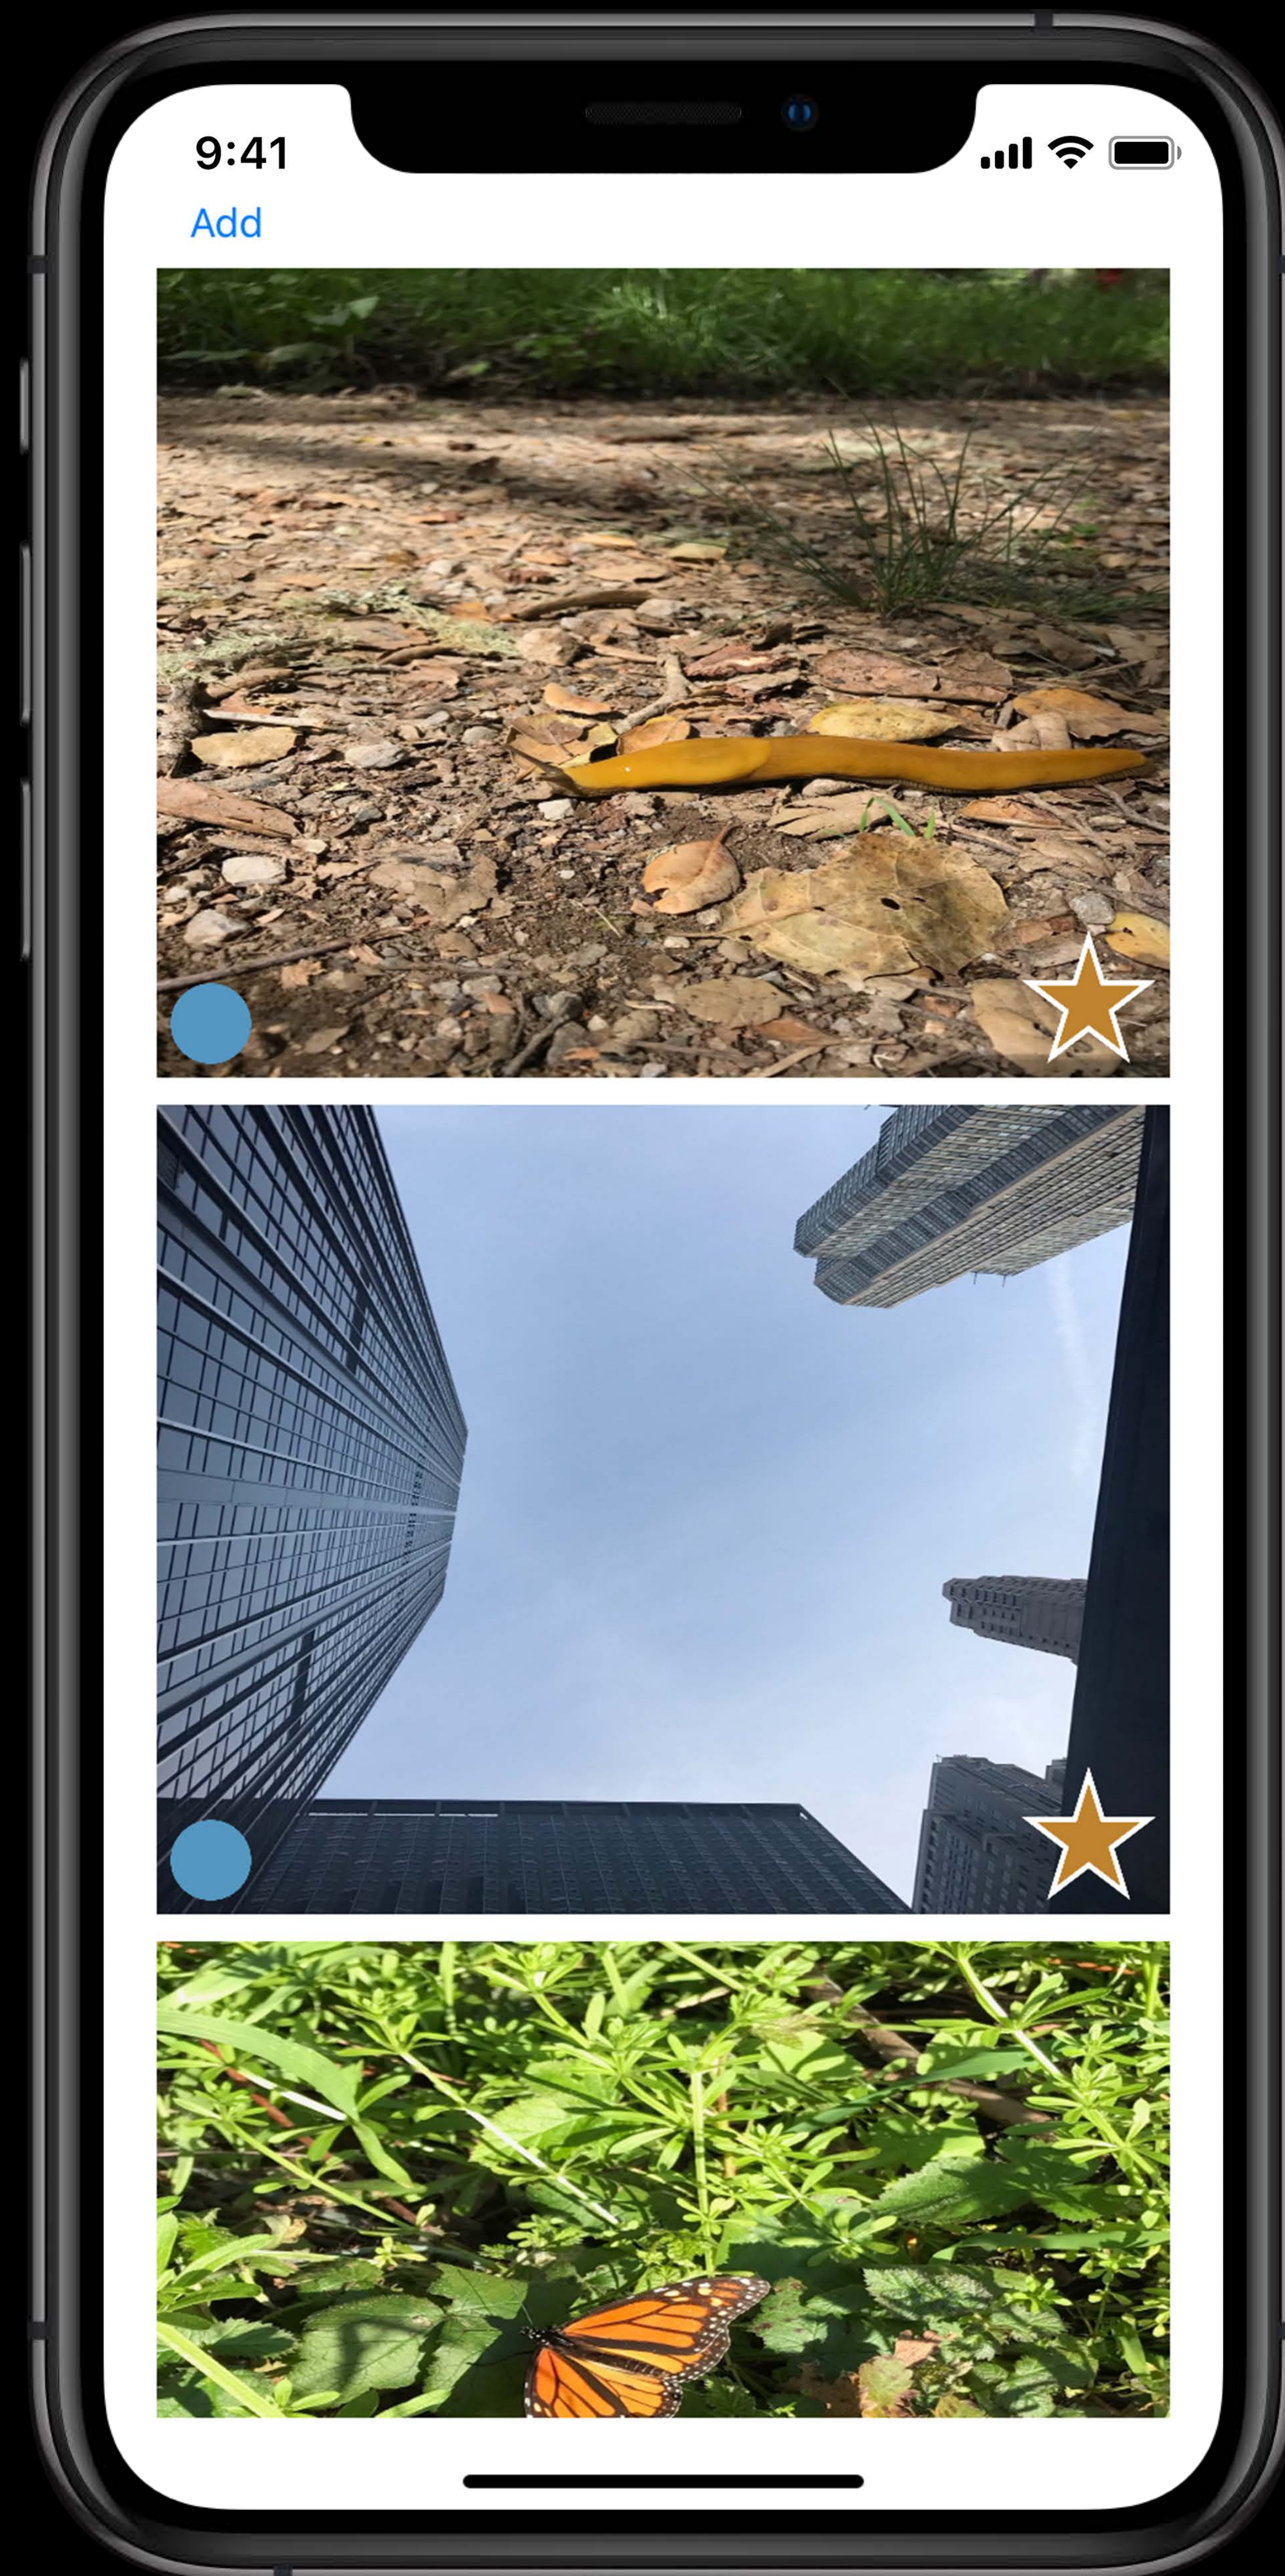


# Image Assets

Growing in size

JPEG Size: 24.6MB

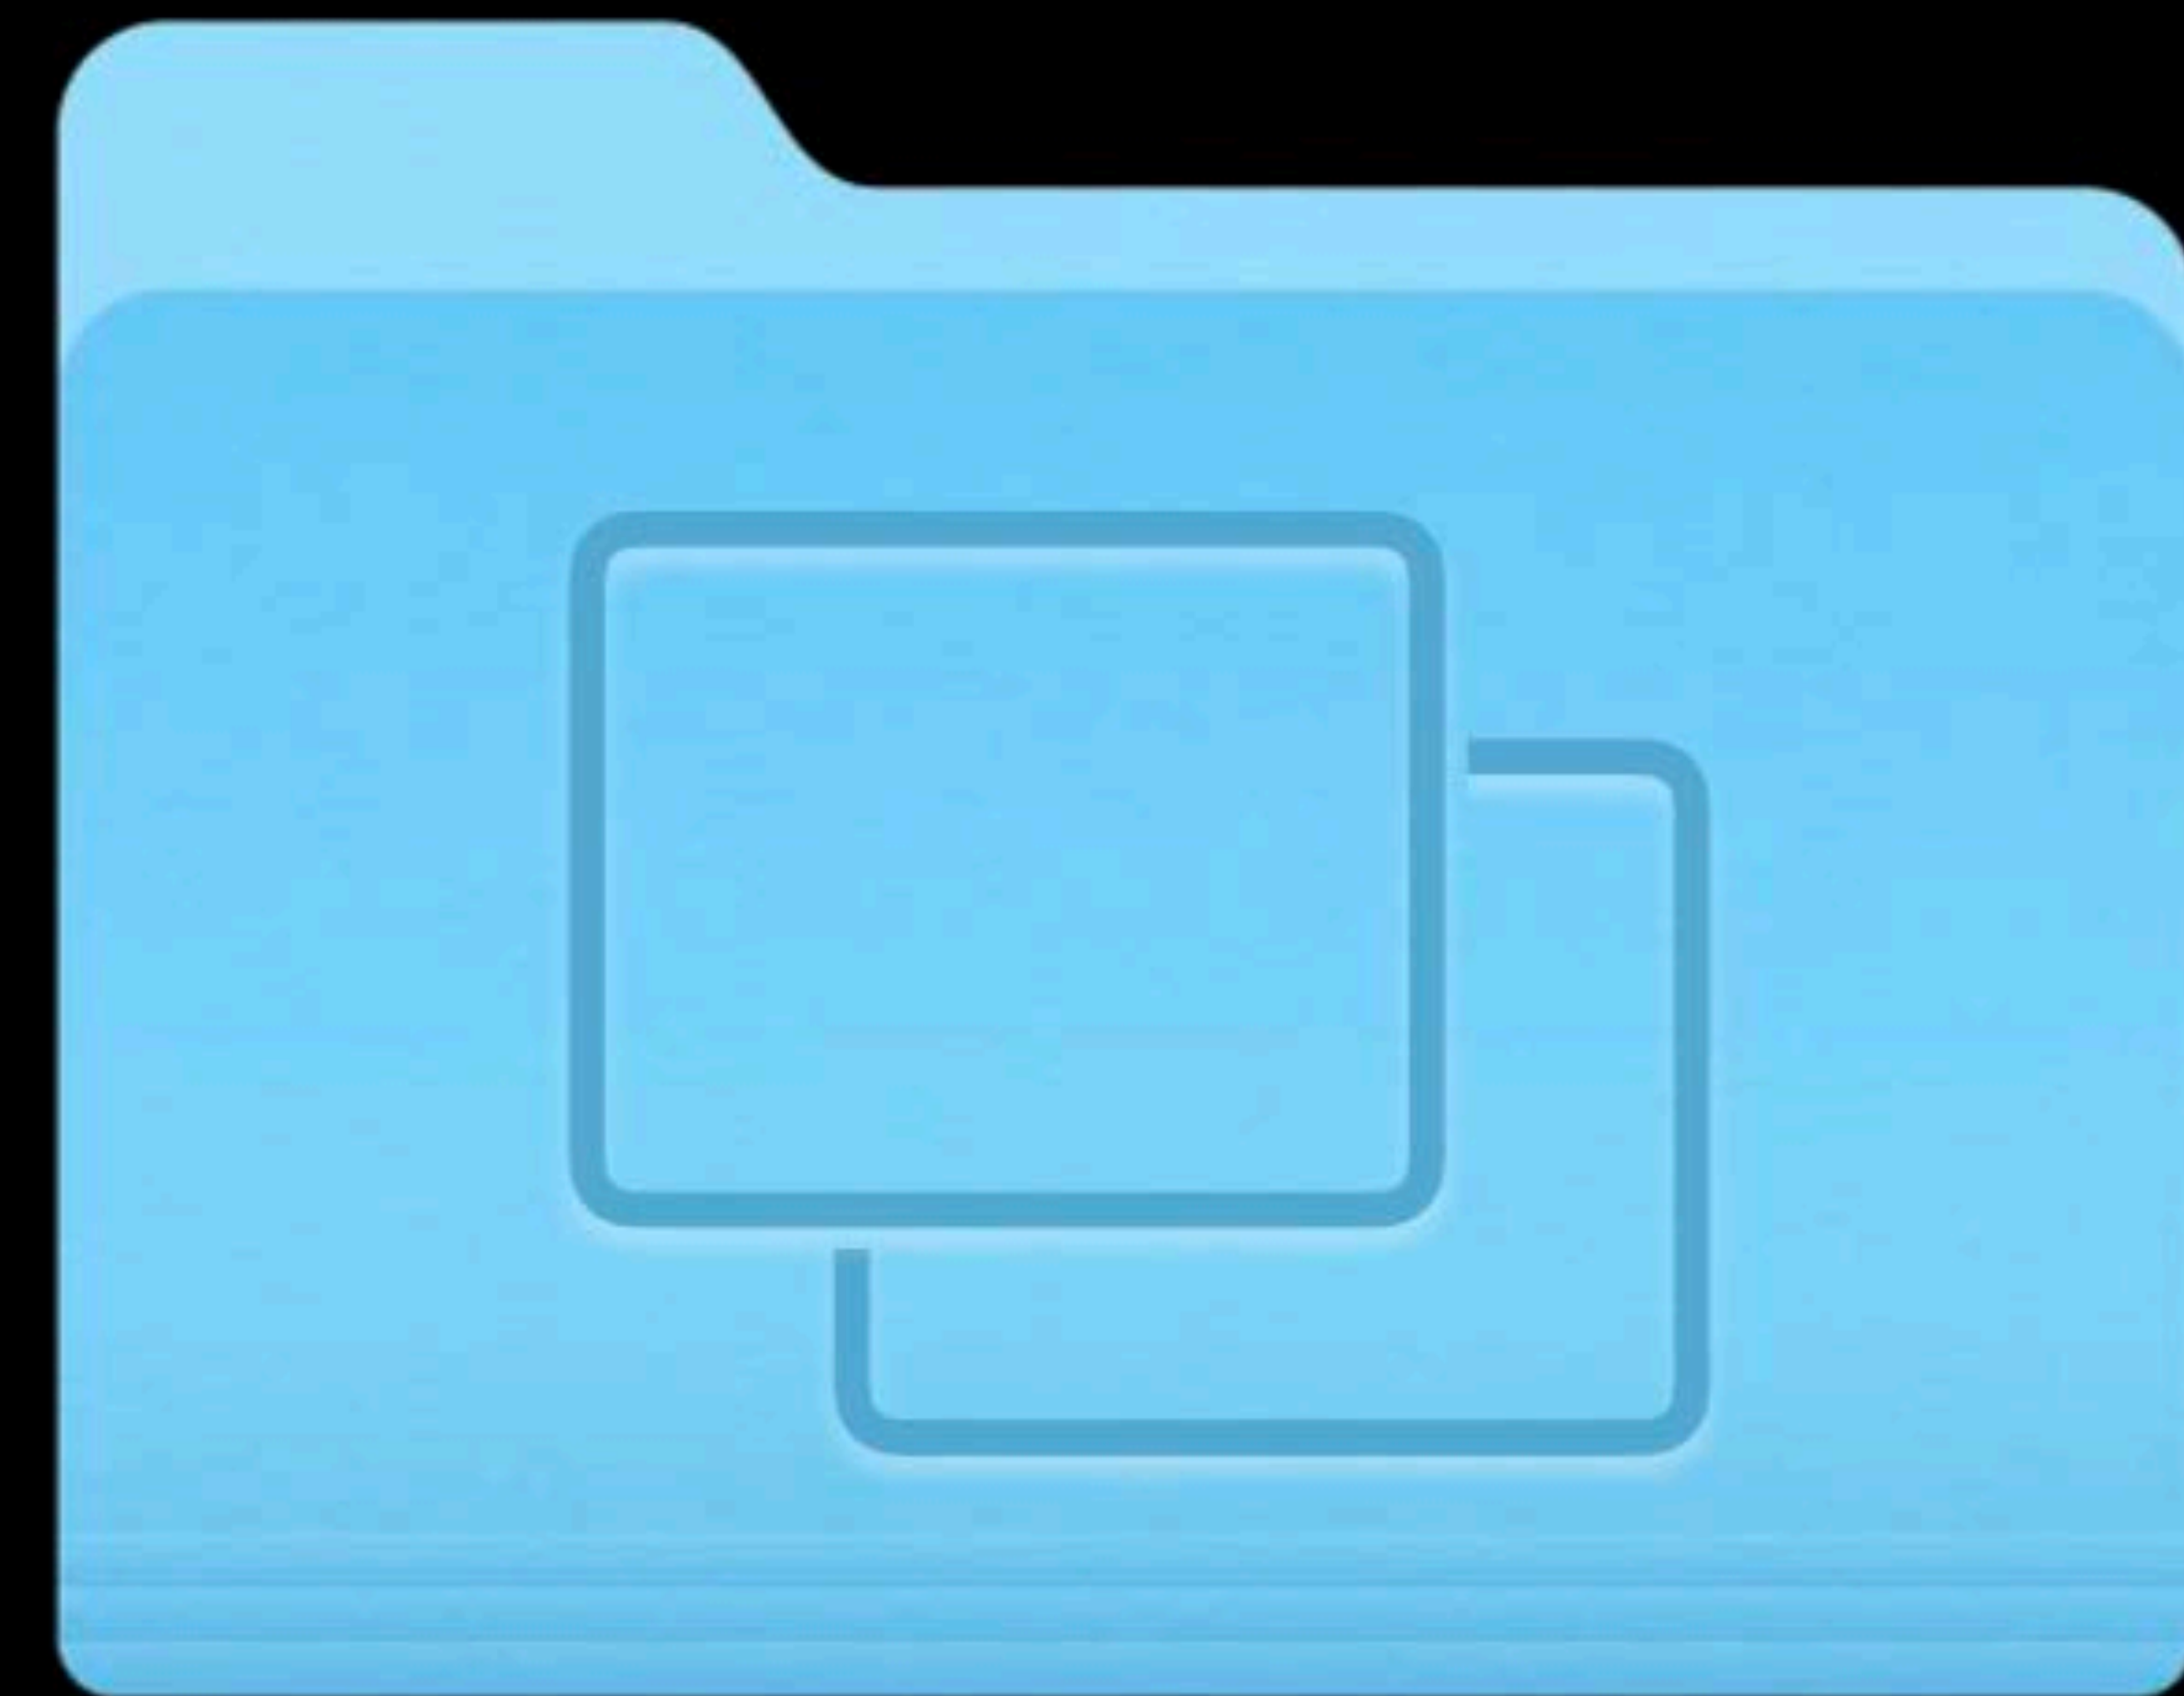
HEIC Size: 17.9MB





# Asset Catalog

Simple app resource management

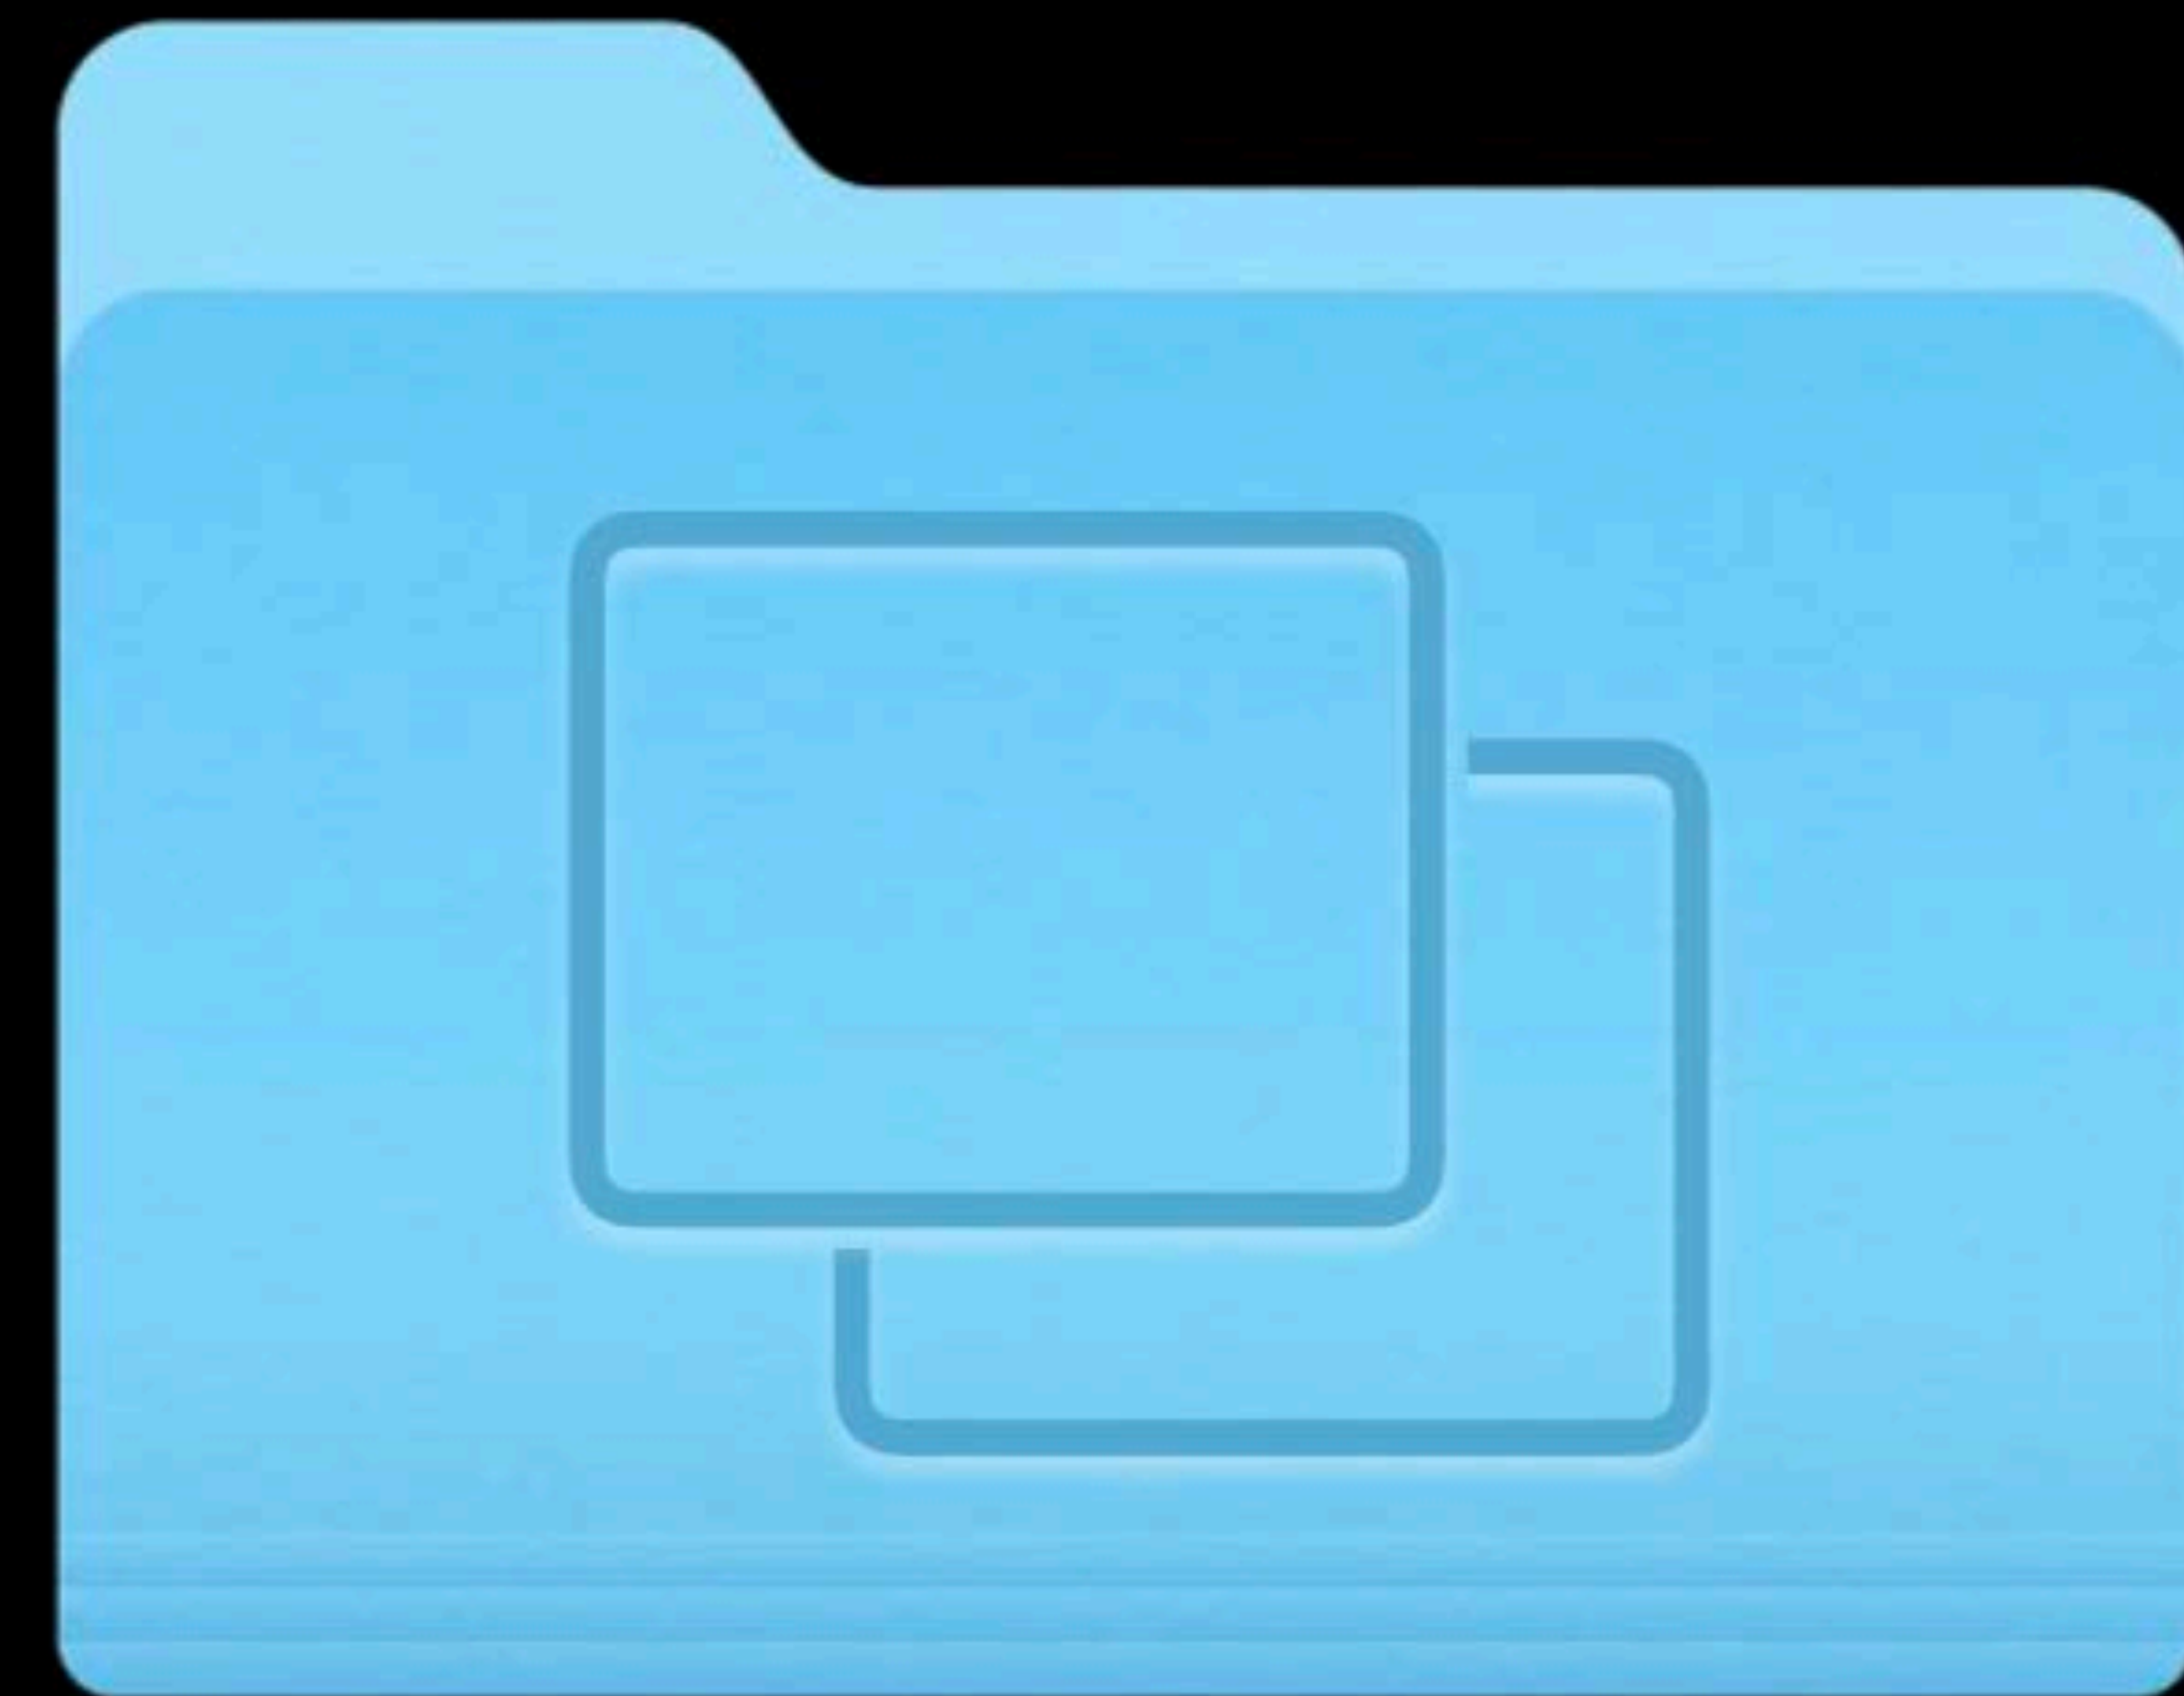




# Asset Catalog

Simple app resource management

App icon

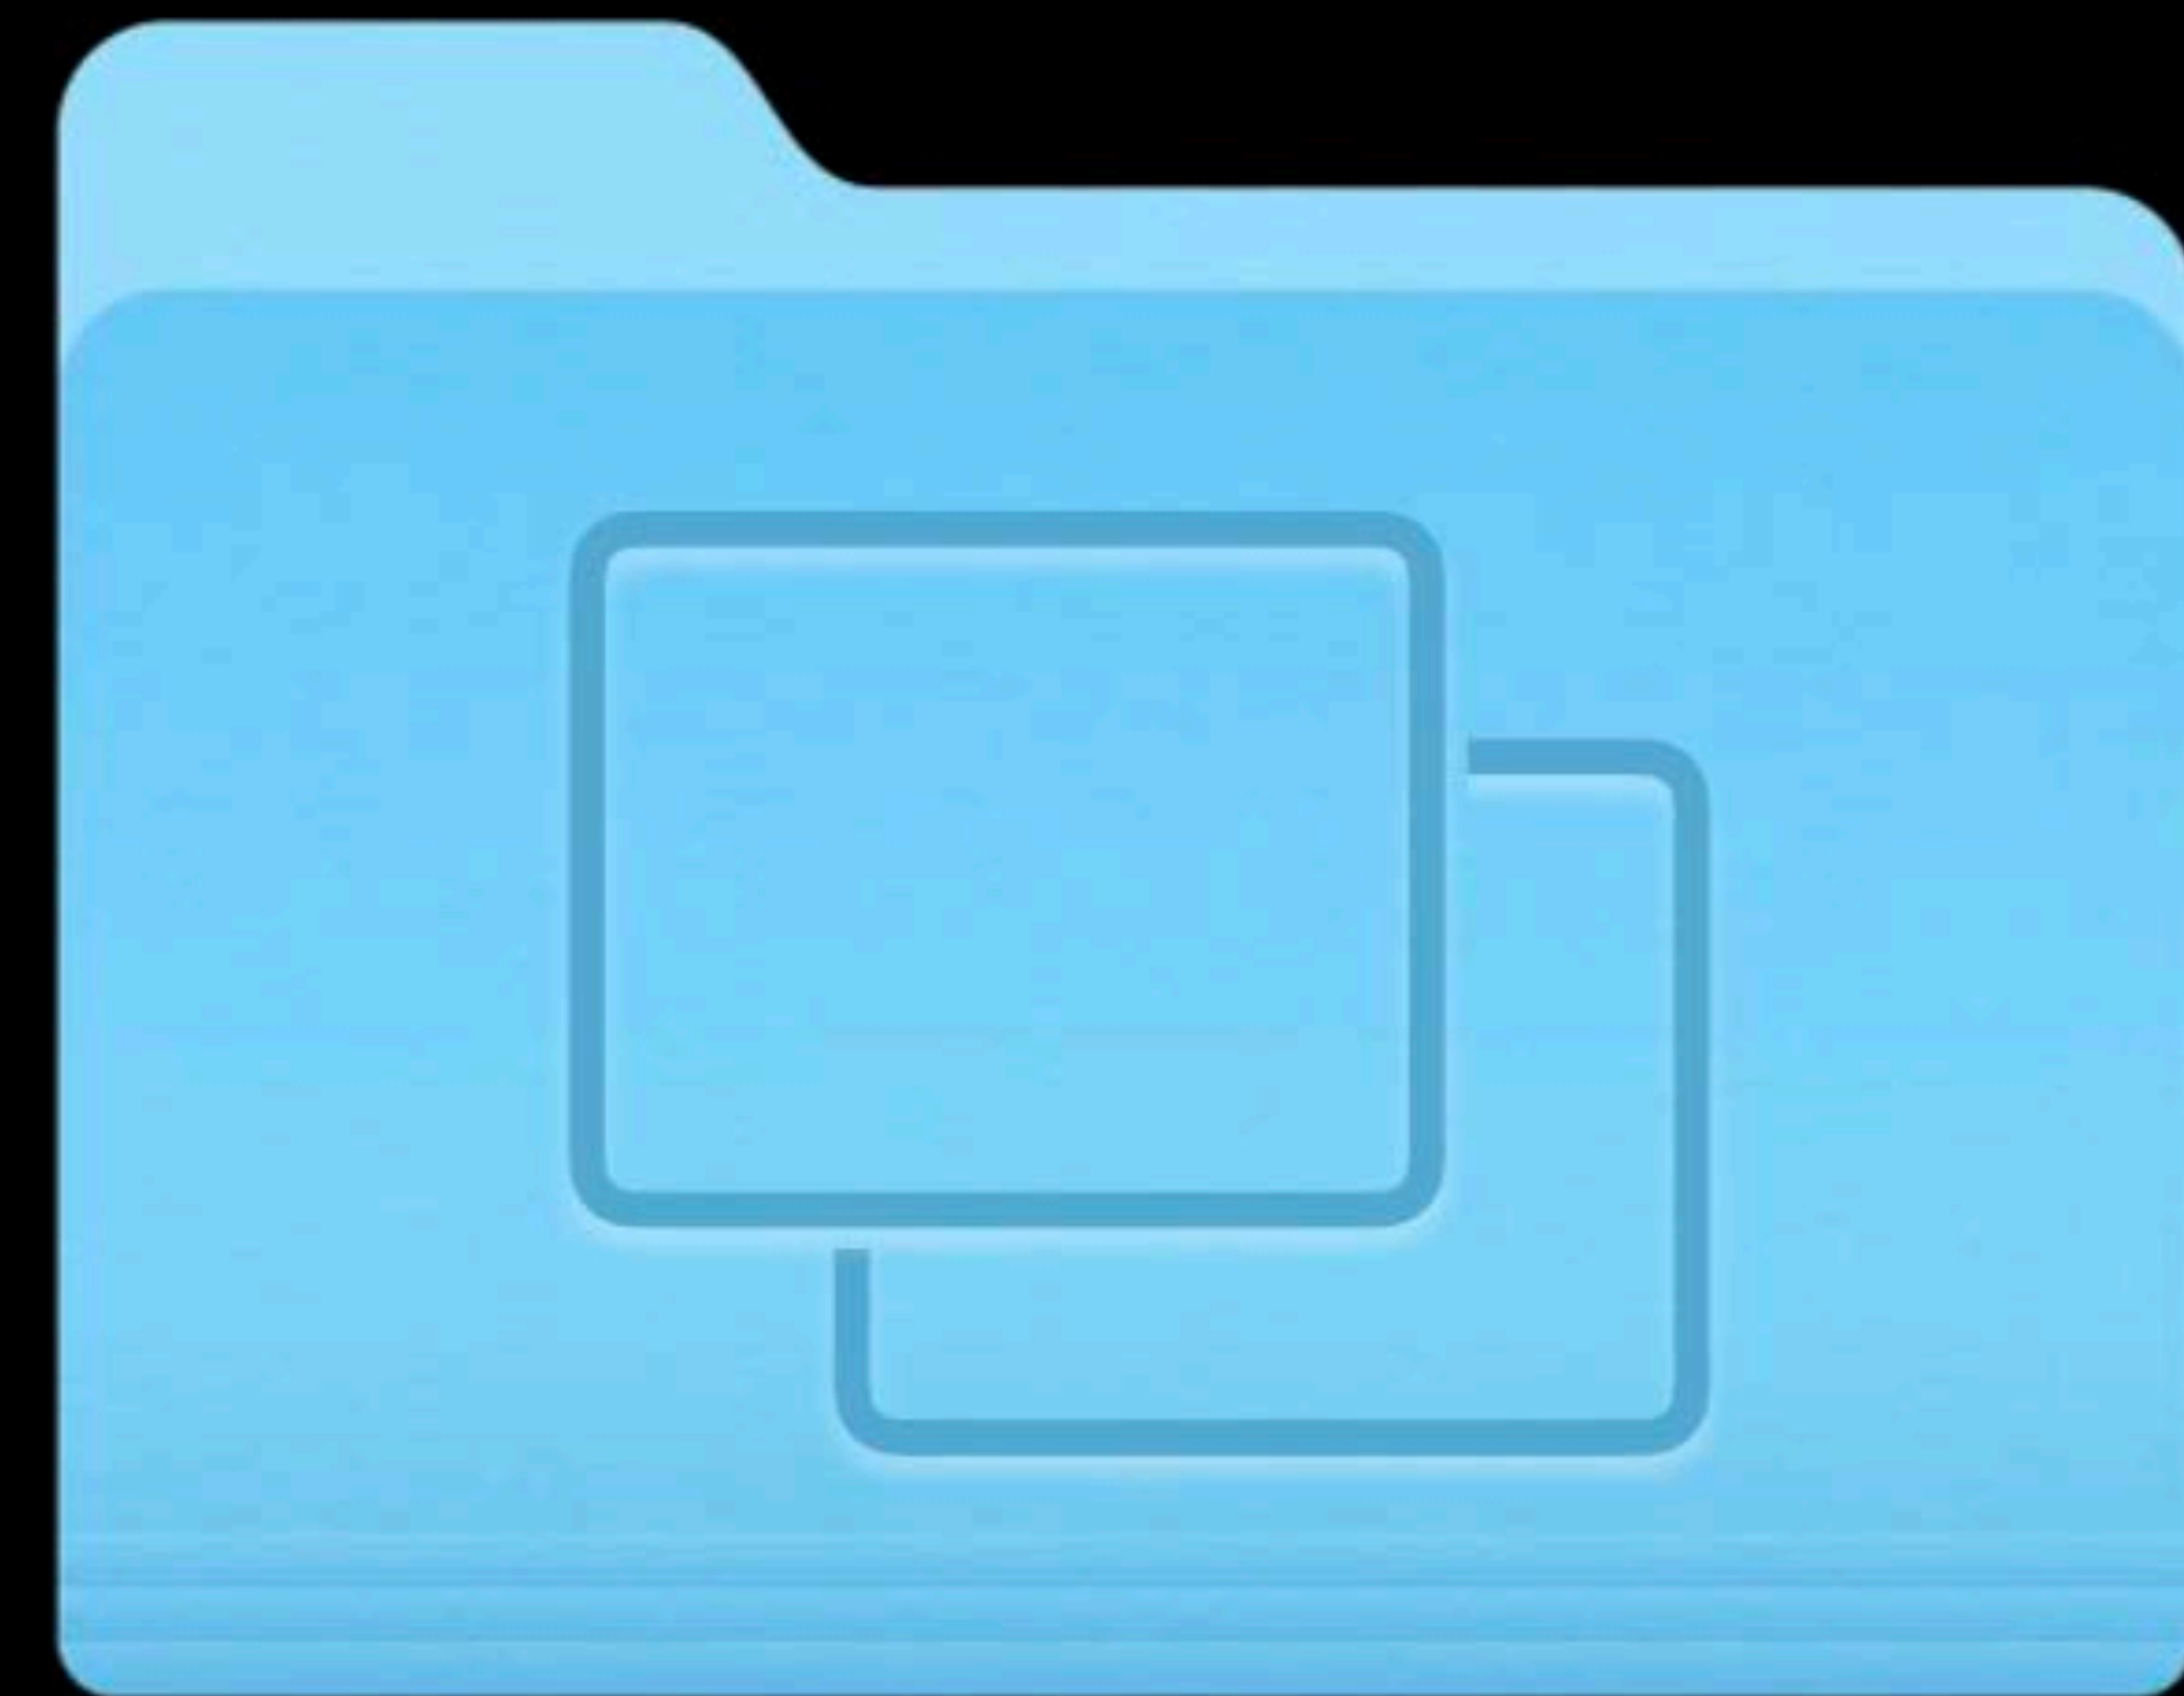


# Asset Catalog

Simple app resource management

App icon

Device and scale variants



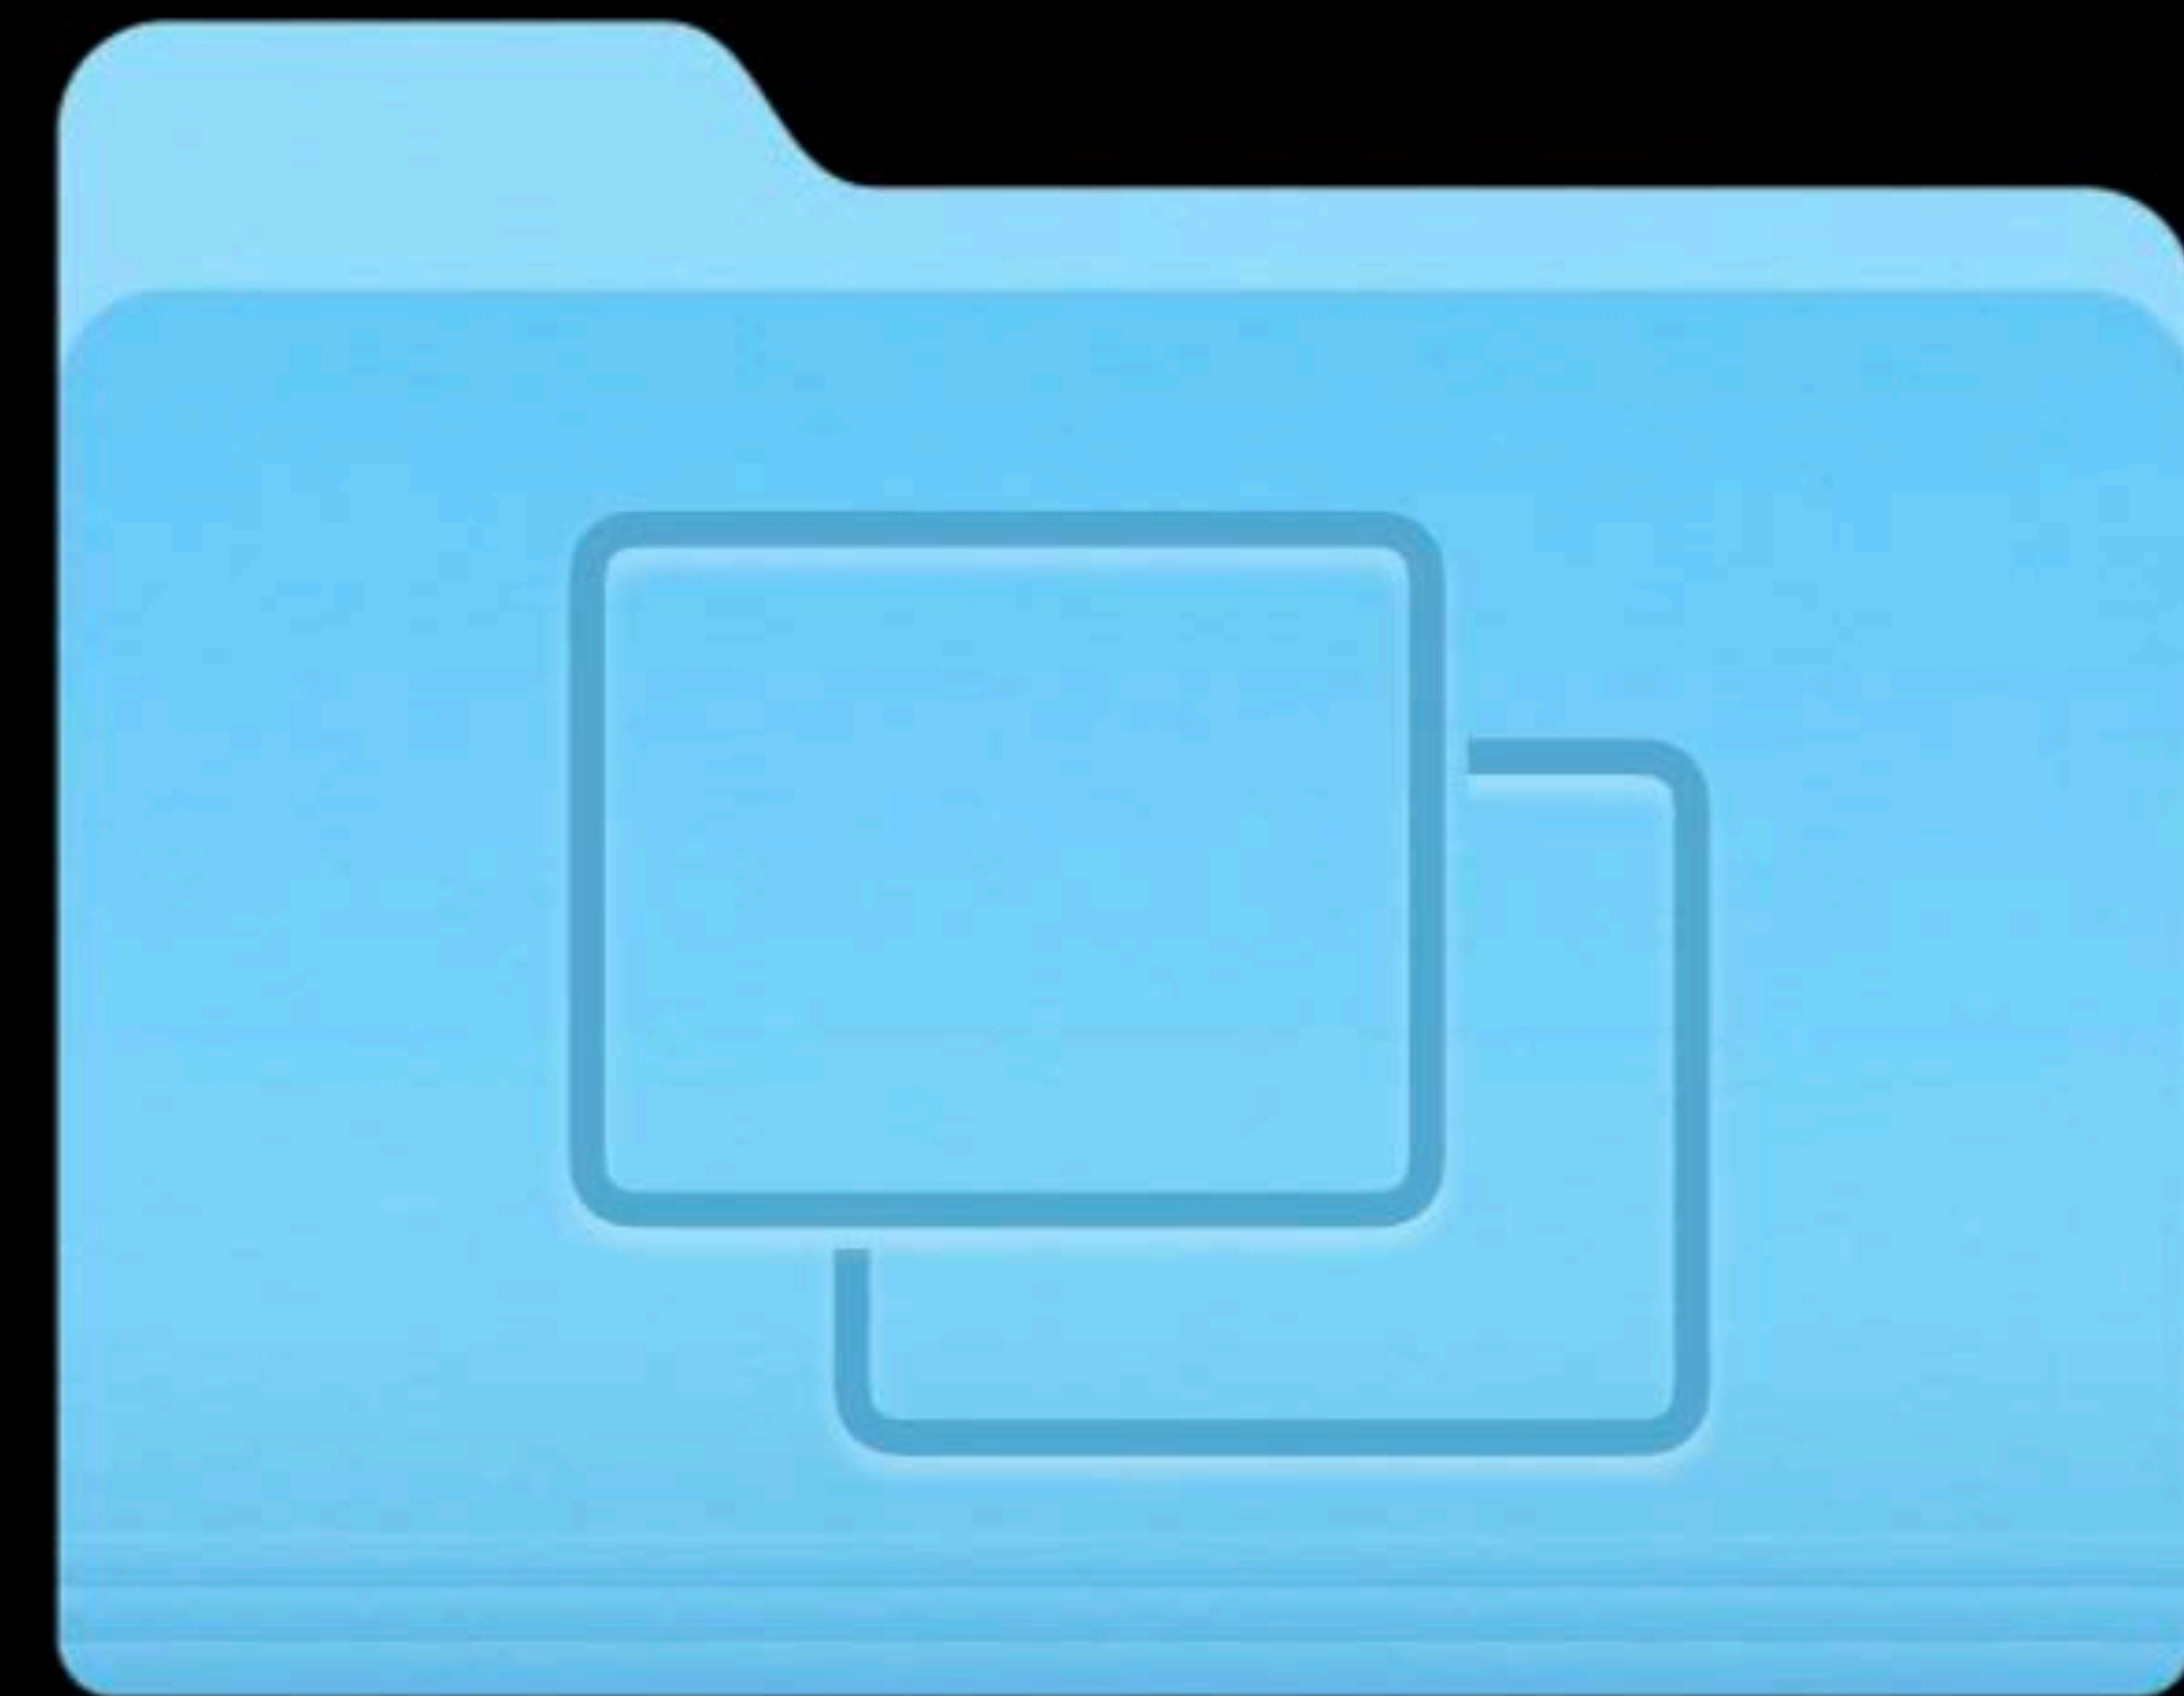
# Asset Catalog

Simple app resource management

App icon

Device and scale variants

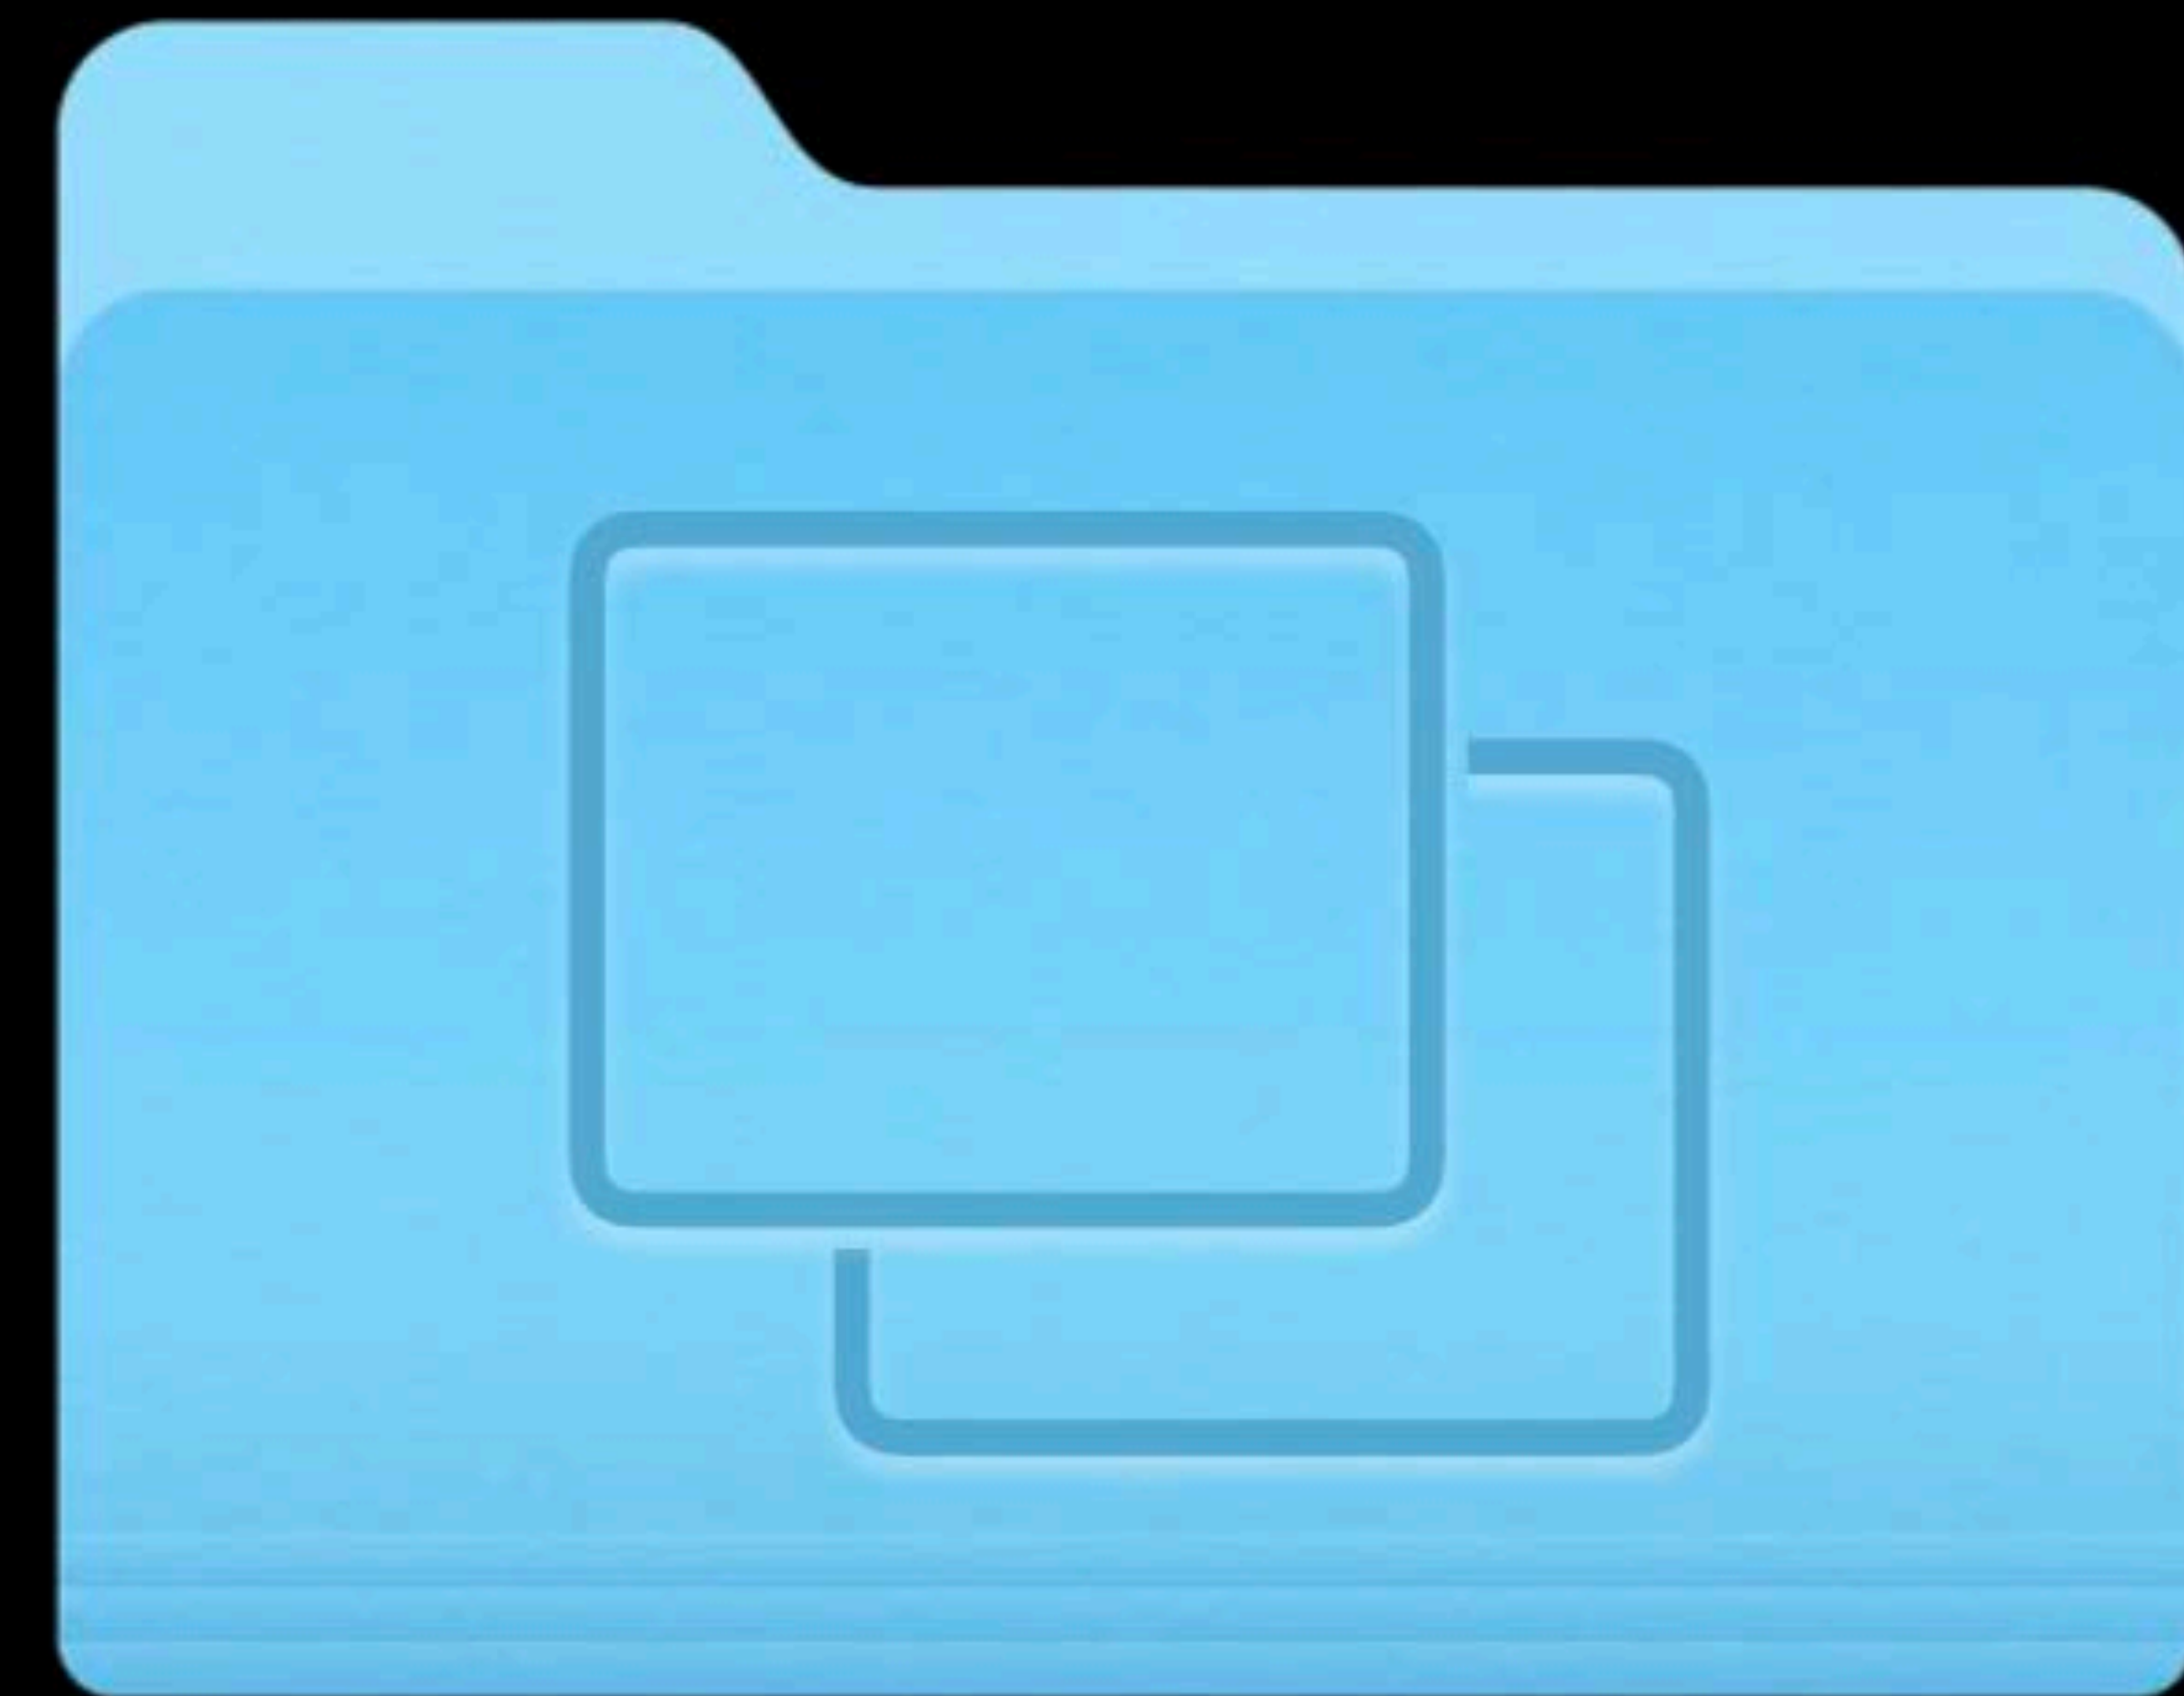
On demand resources



# Asset Catalog

Storage and performance benefits

Storage efficiency



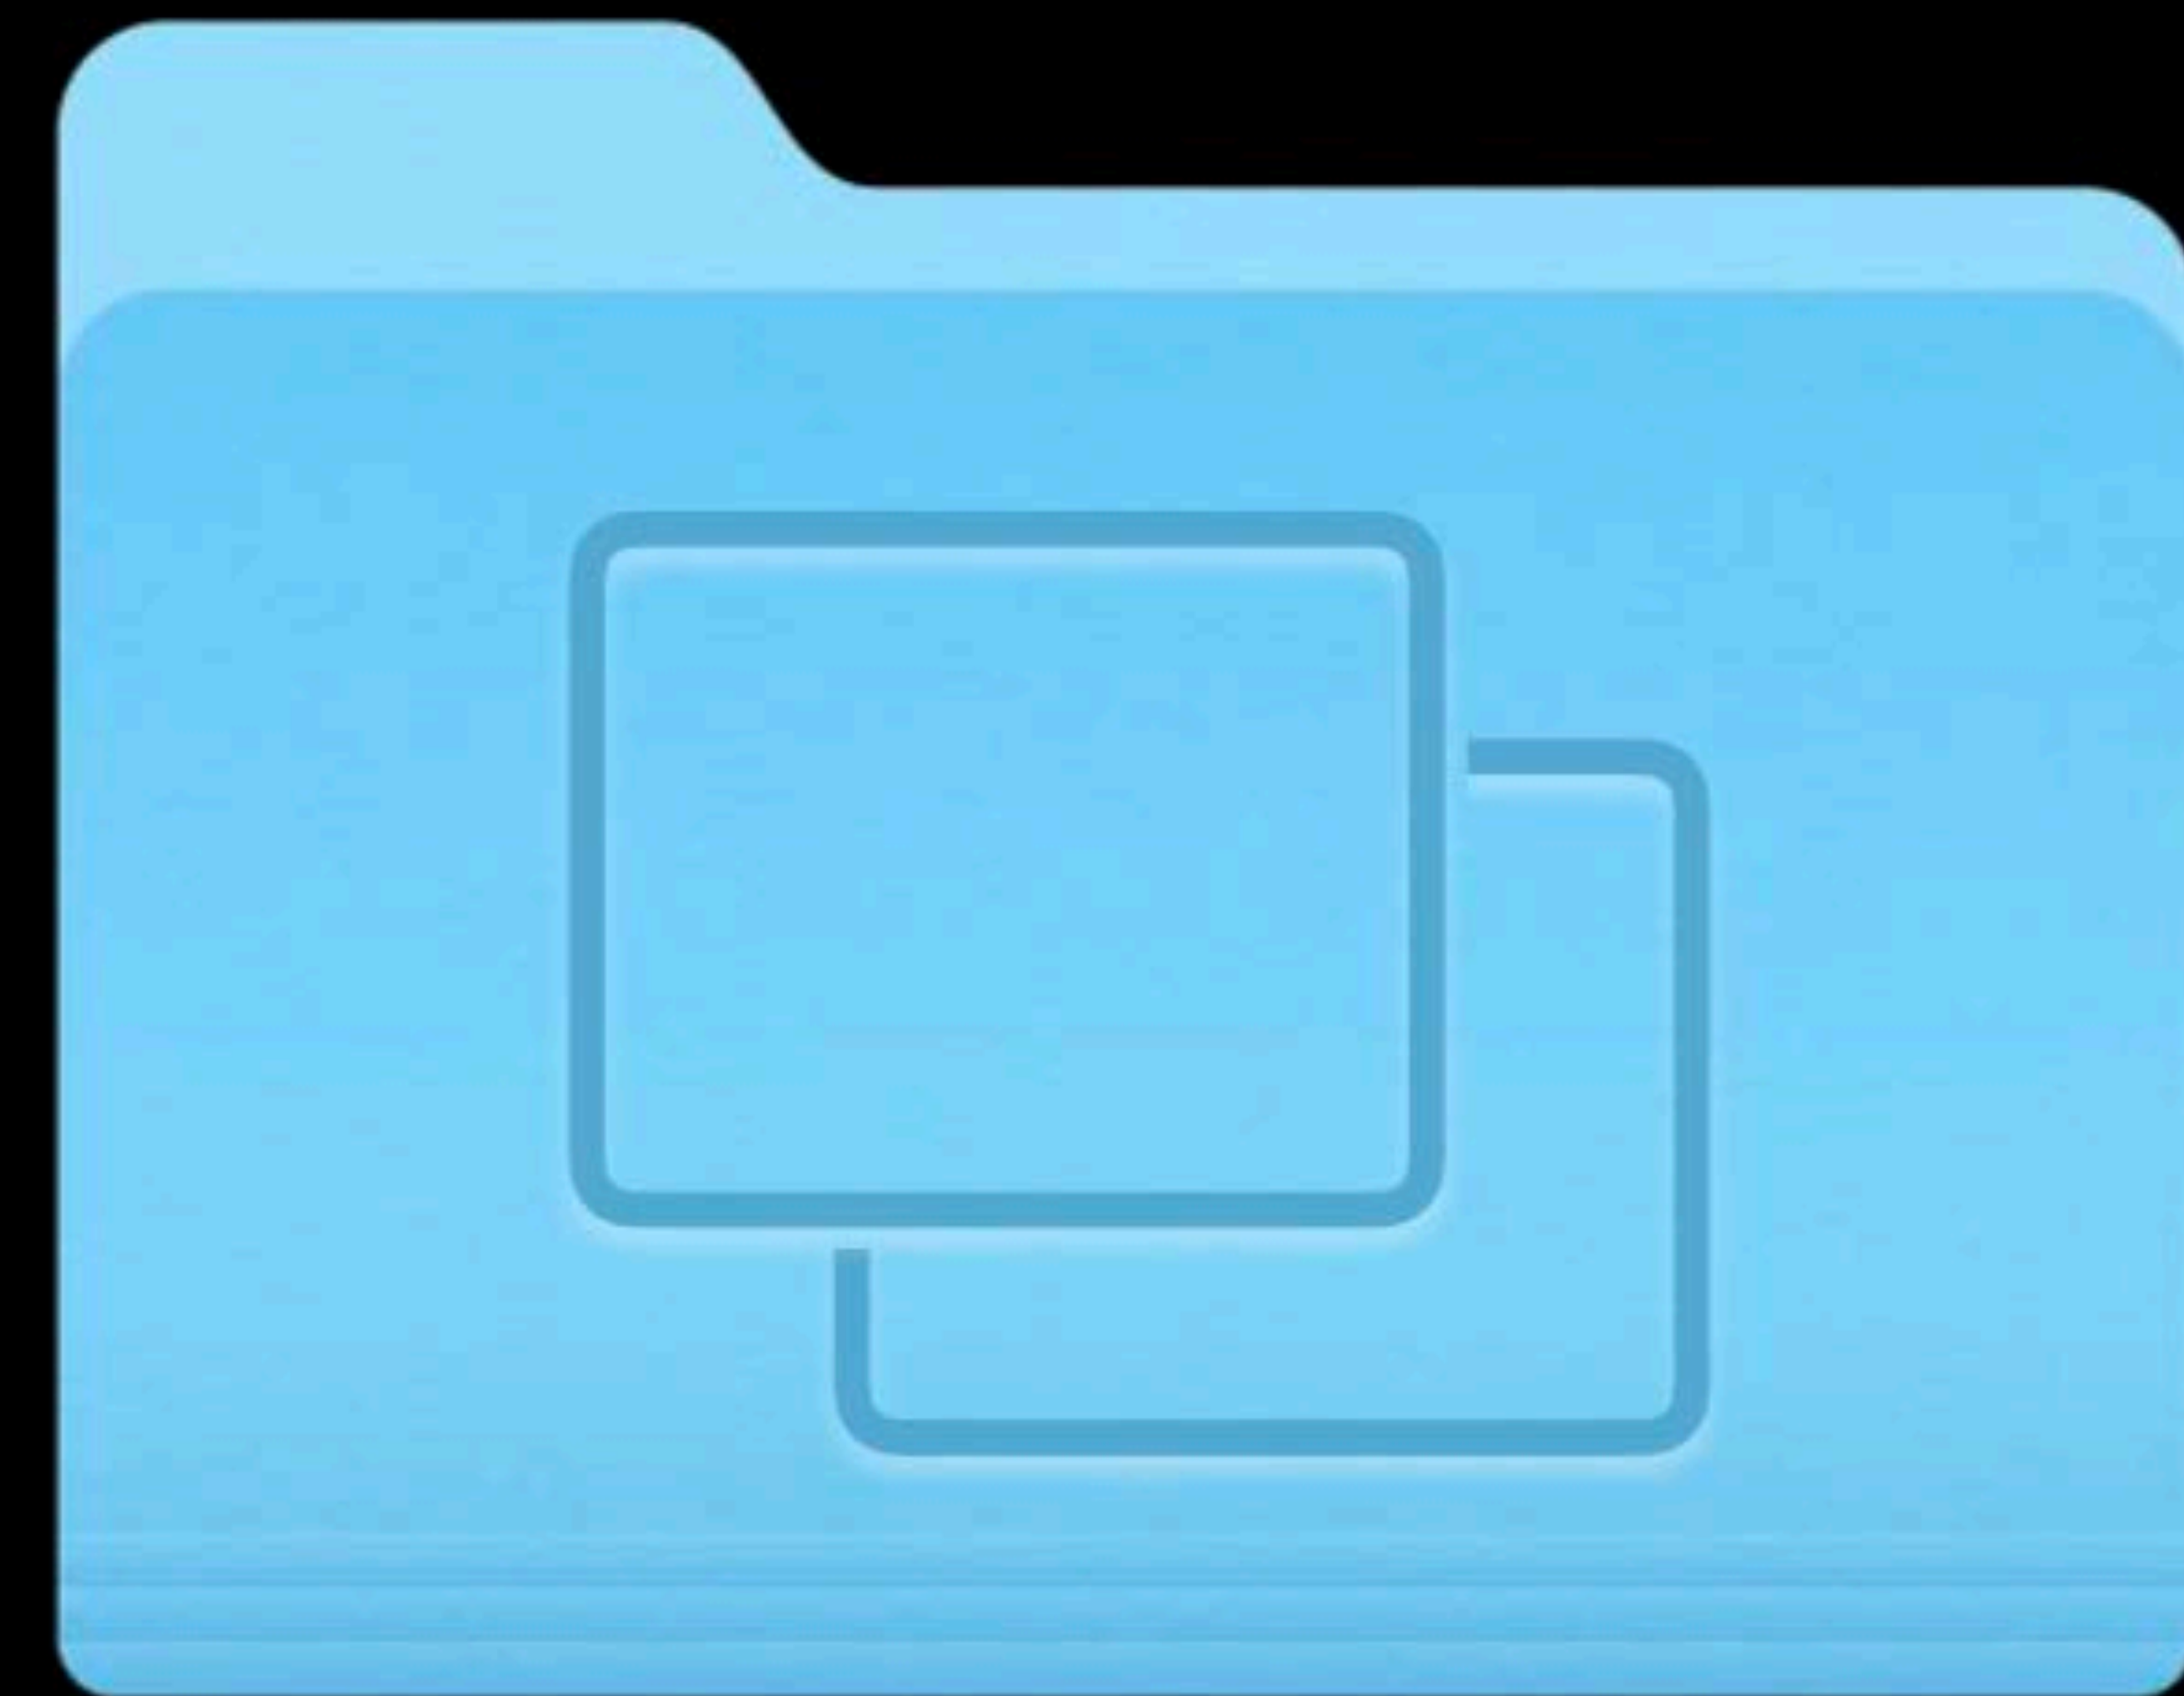


# Asset Catalog

Storage and performance benefits

Storage efficiency

- On-disk foot print

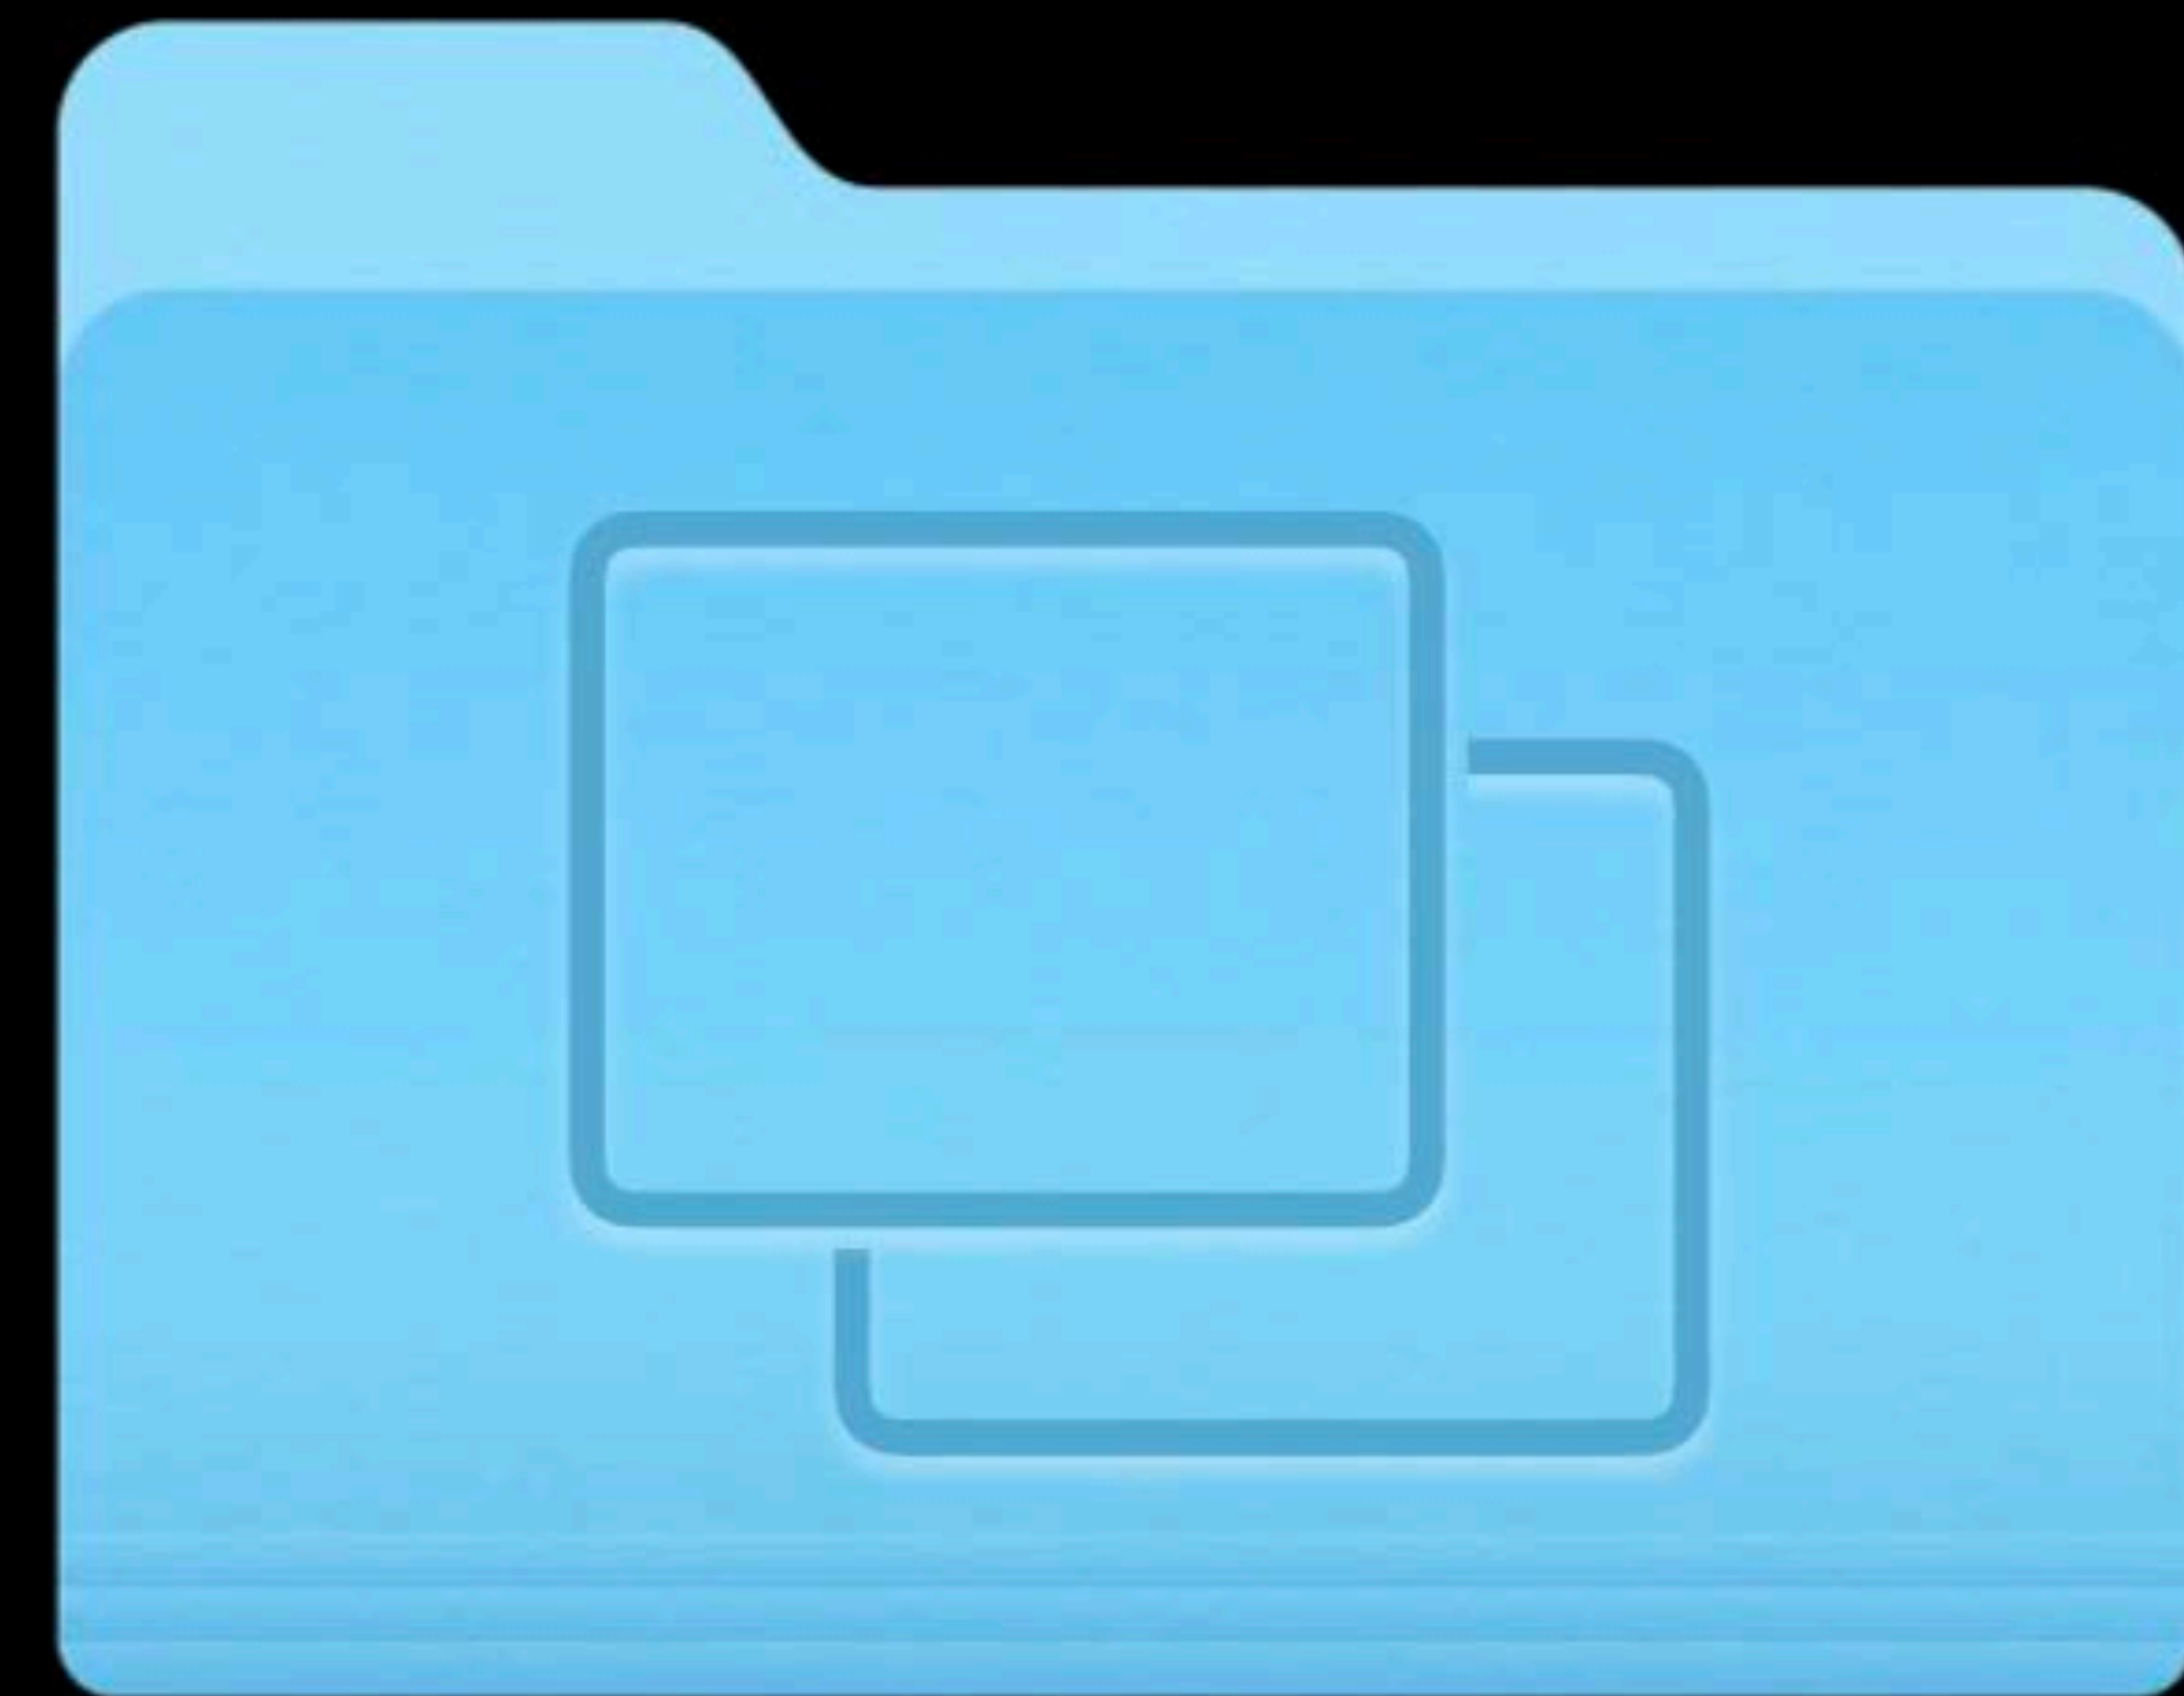


# Asset Catalog

Storage and performance benefits

Storage efficiency

- On-disk foot print
- App slicing (iOS)



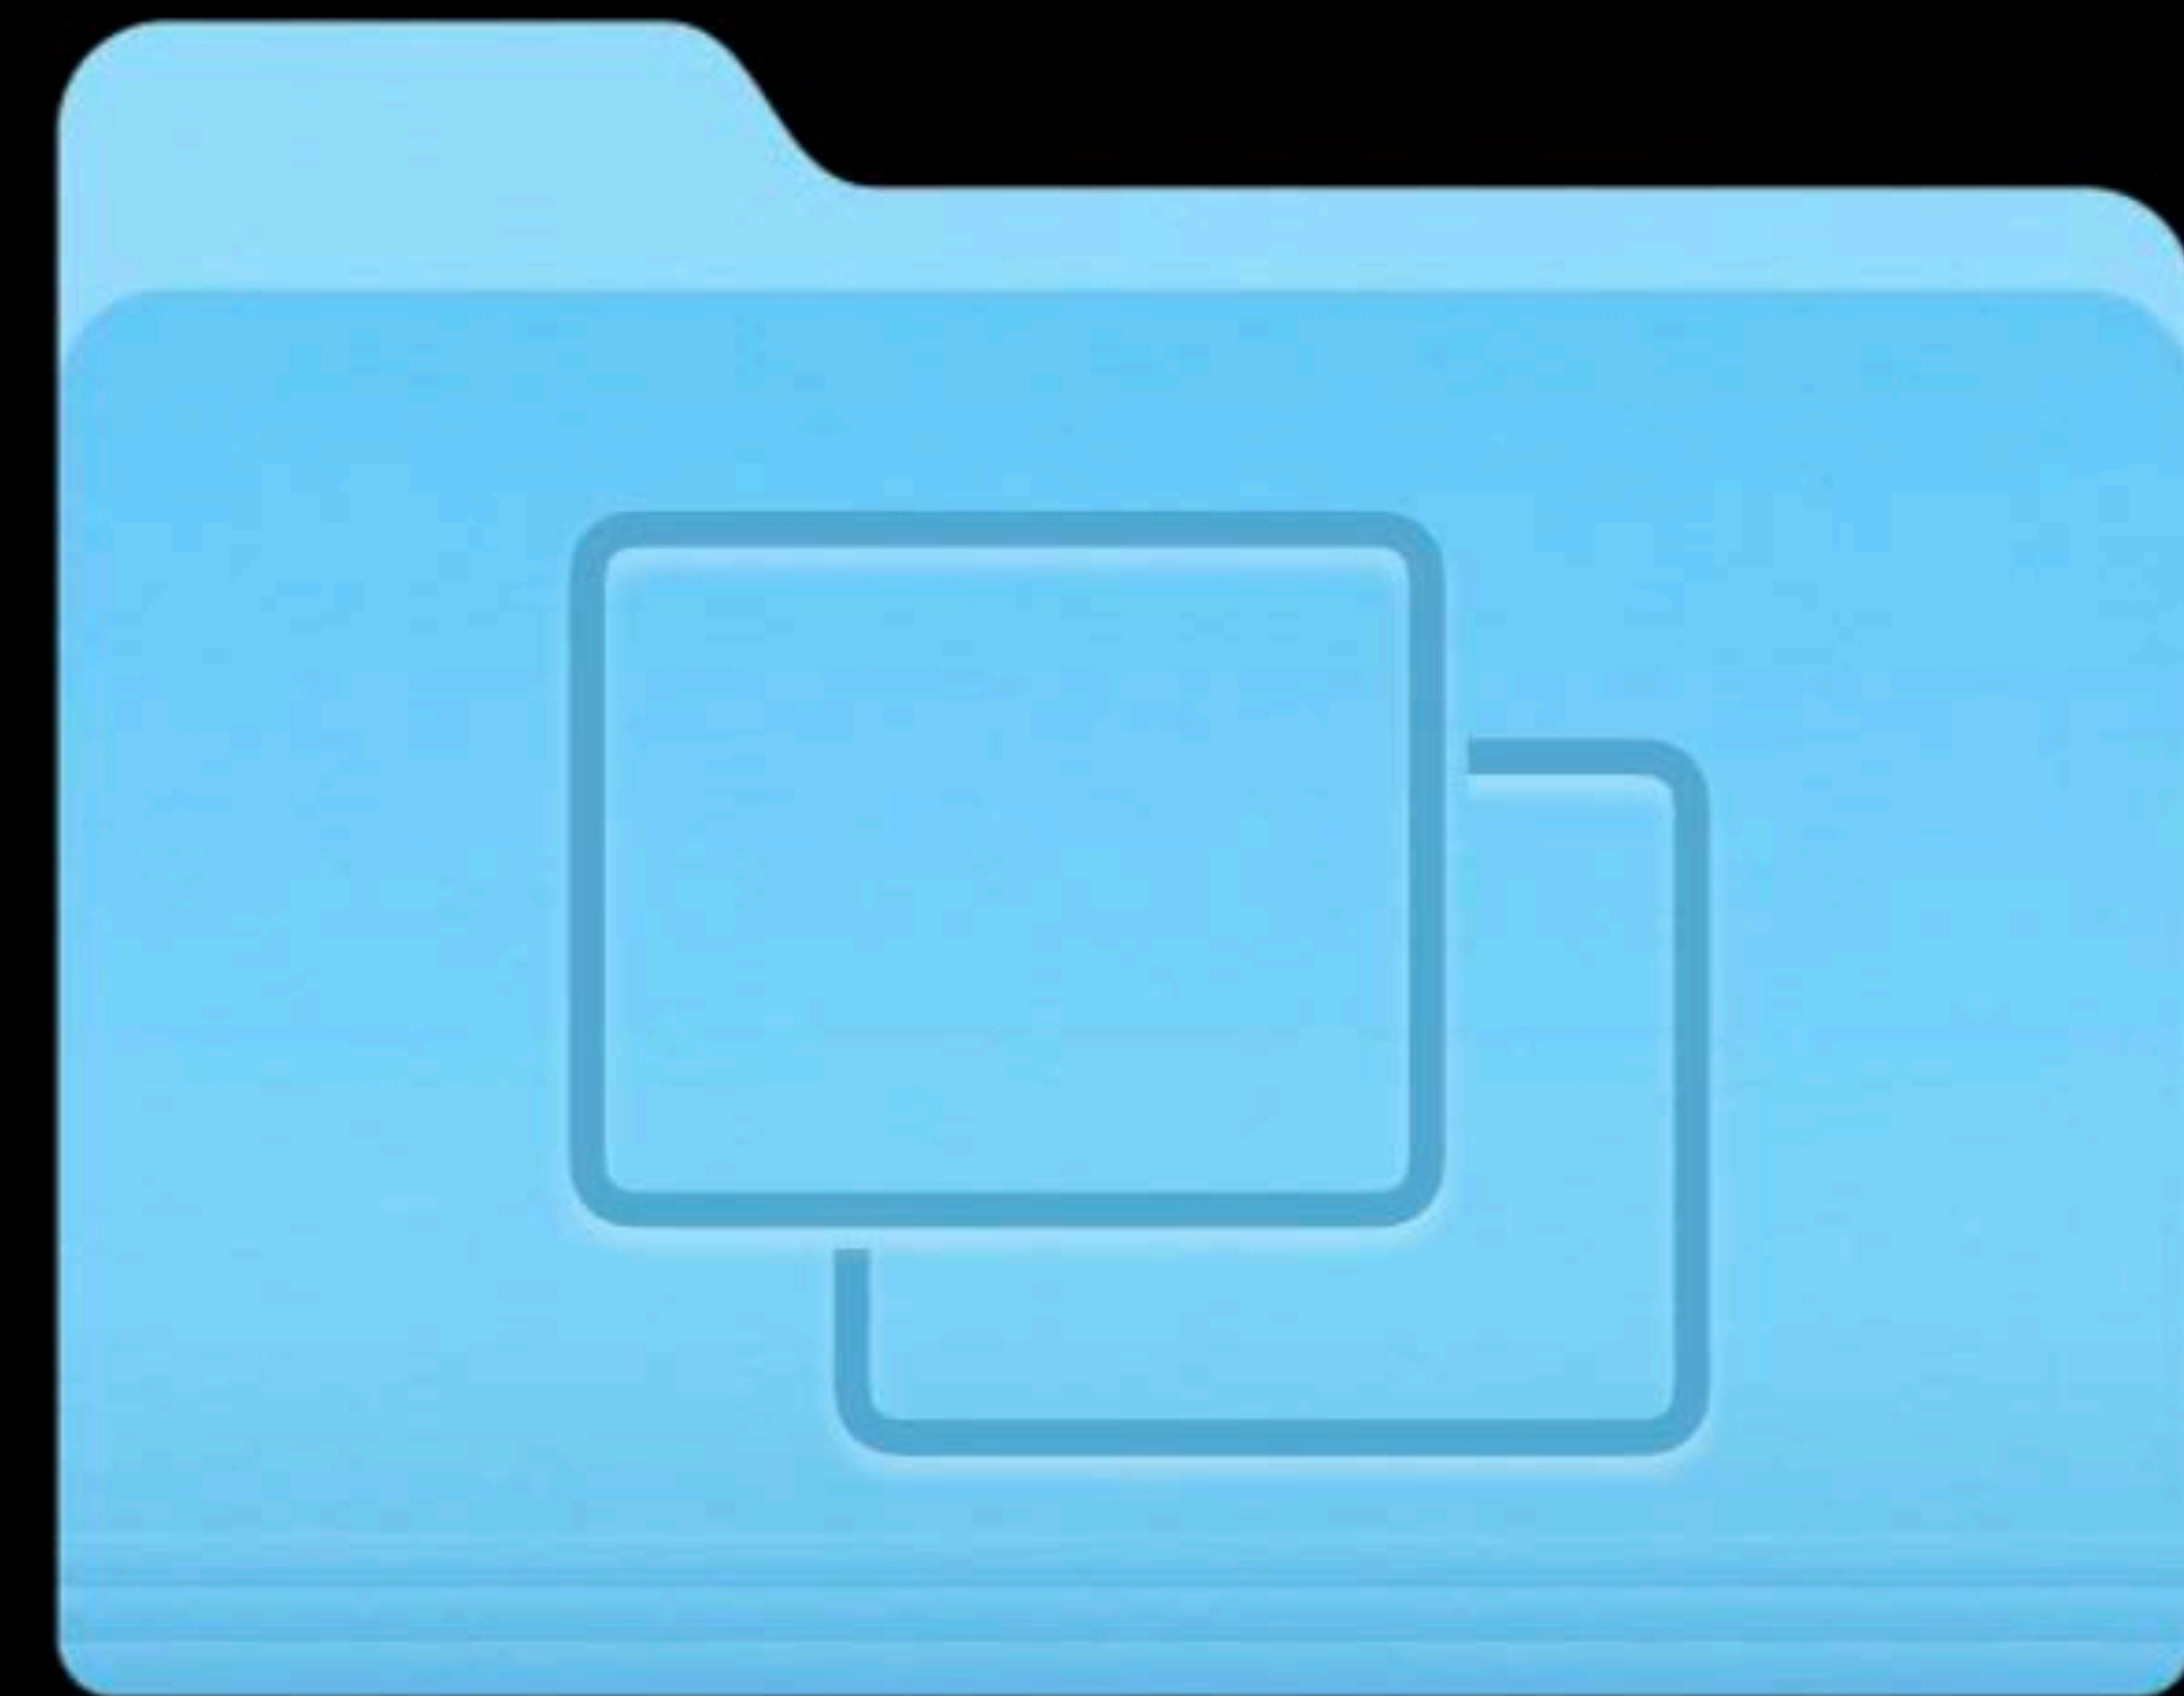
# Asset Catalog

Storage and performance benefits

Storage efficiency

- On-disk foot print
- App slicing (iOS)

Performance



# Asset Catalog

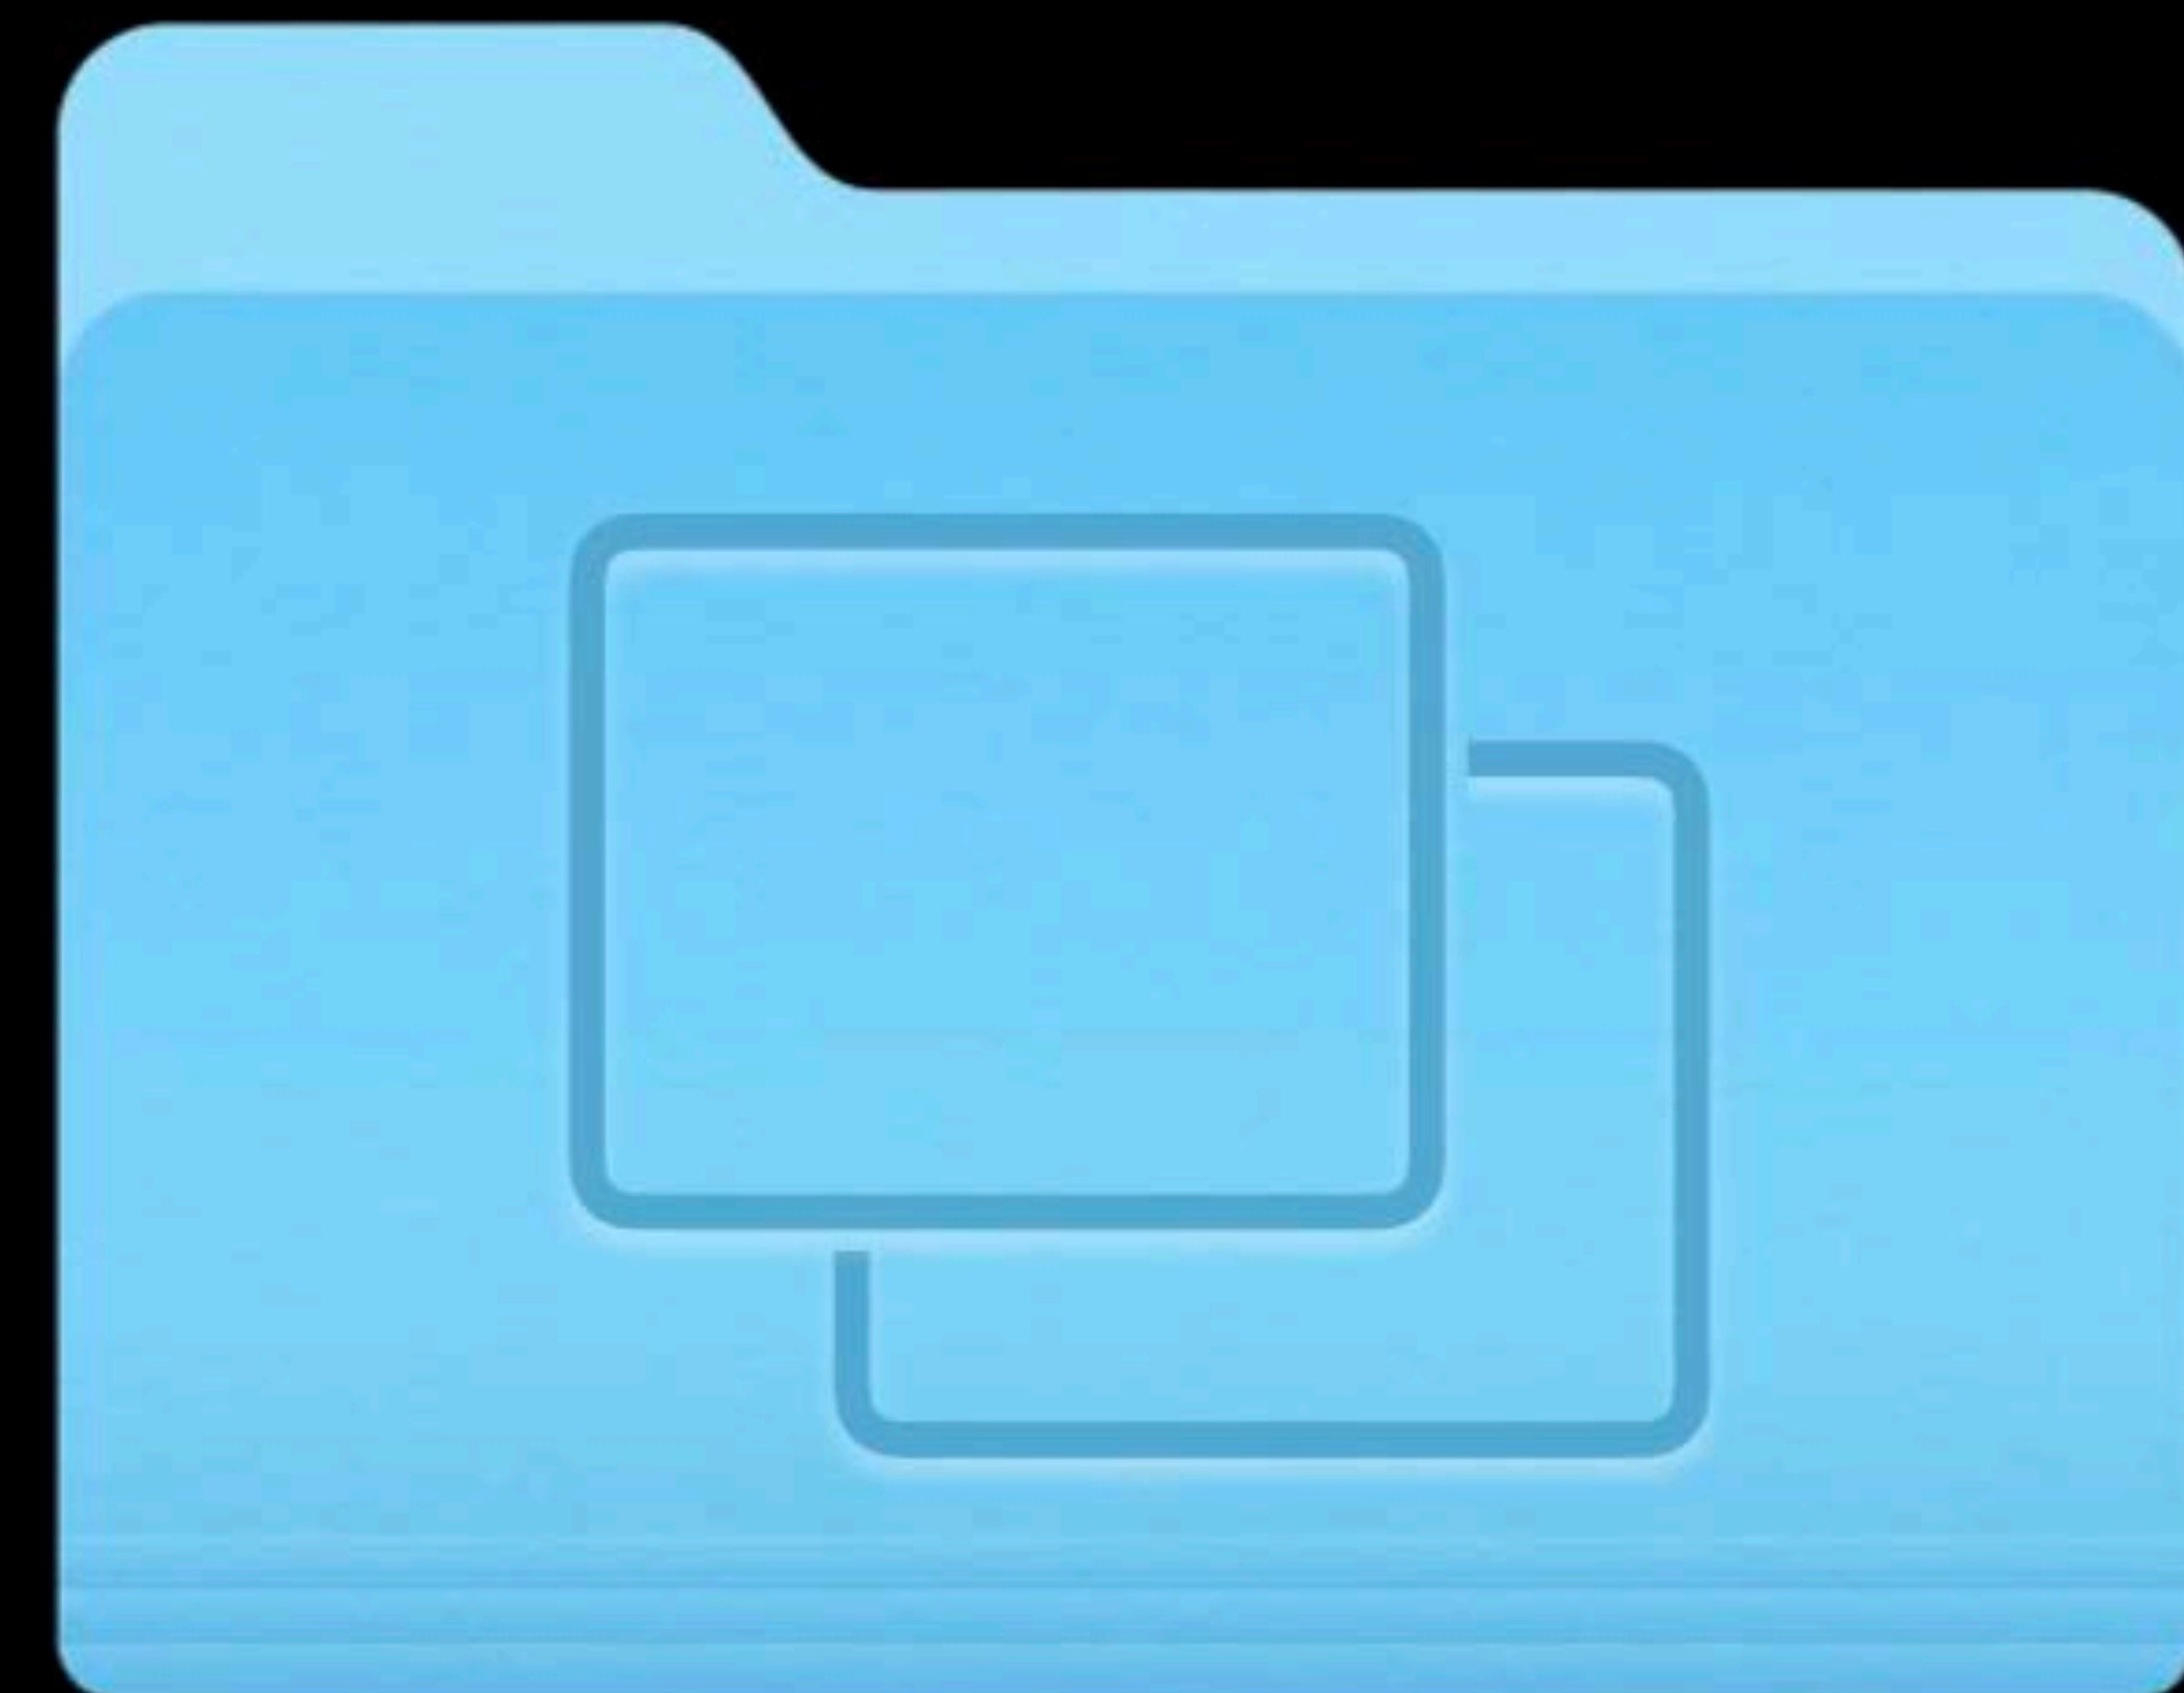
Storage and performance benefits

Storage efficiency

- On-disk foot print
- App slicing (iOS)

Performance

- Image loading





# Asset Catalog

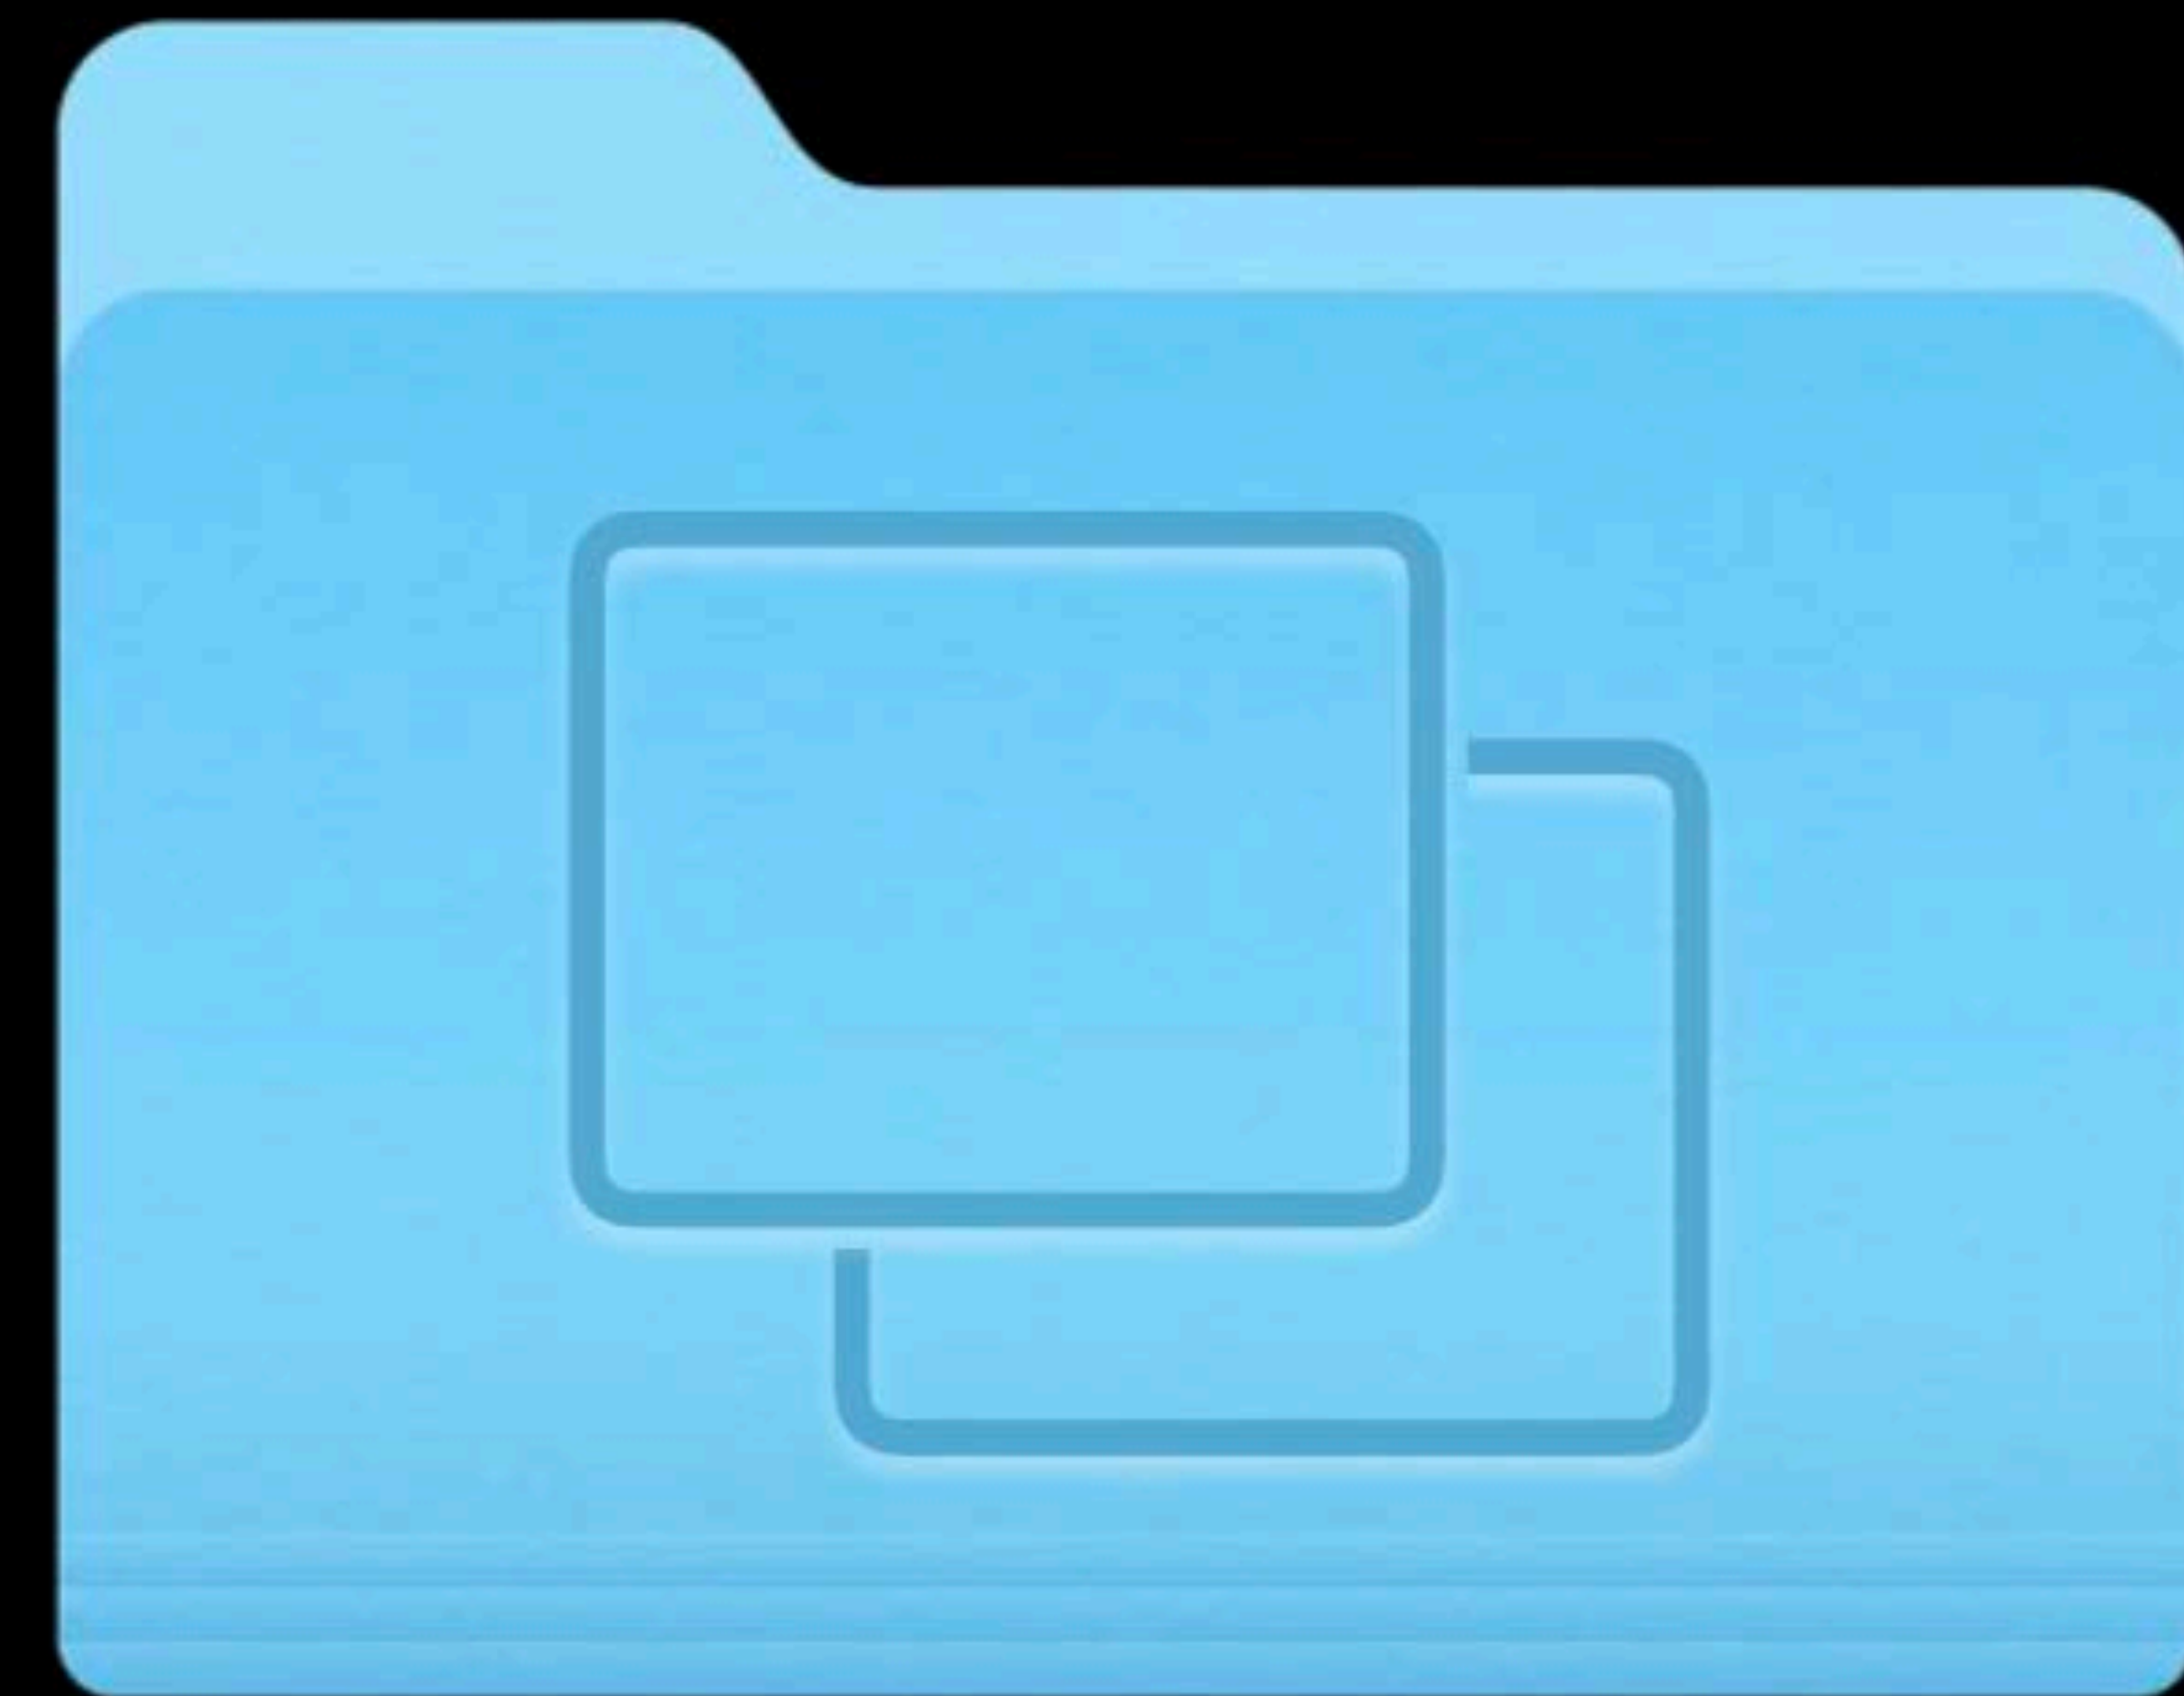
Storage and performance benefits

## Storage efficiency

- On-disk foot print
- App slicing (iOS)

## Performance

- Image loading
- App launch



# Asset Catalog

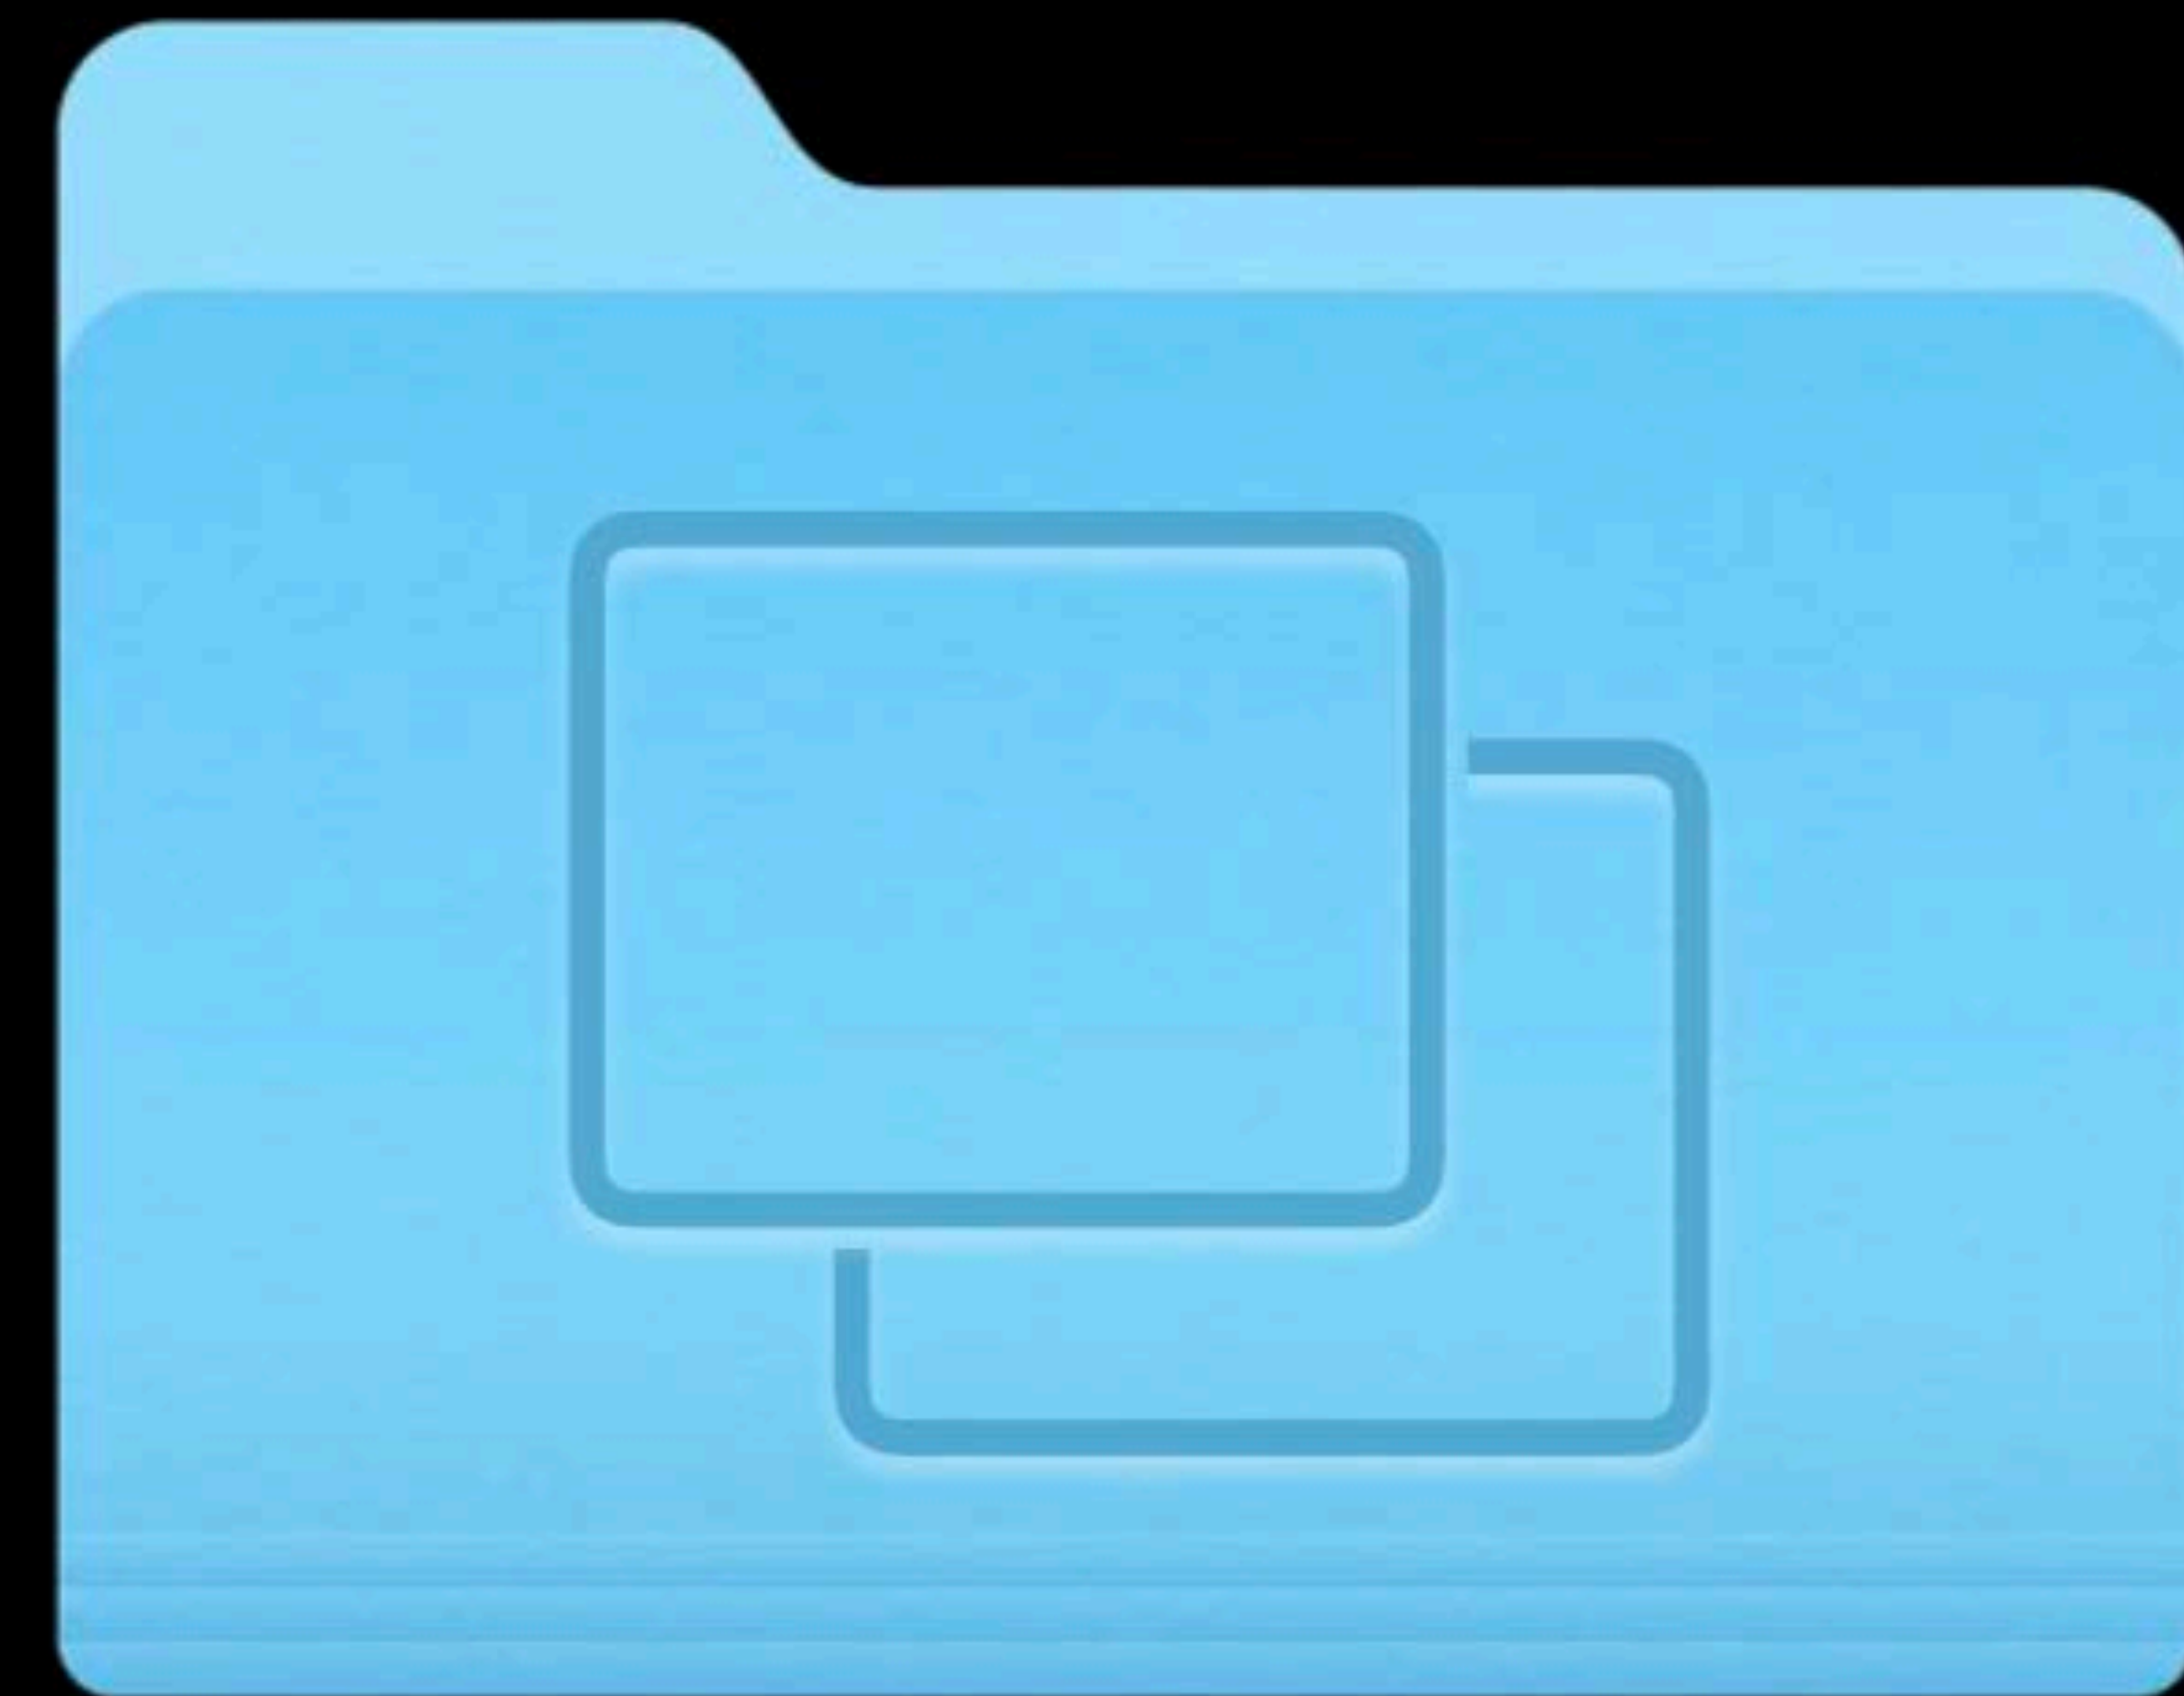
Storage and performance benefits

## Storage efficiency

- On-disk foot print
- App slicing (iOS)

## Performance

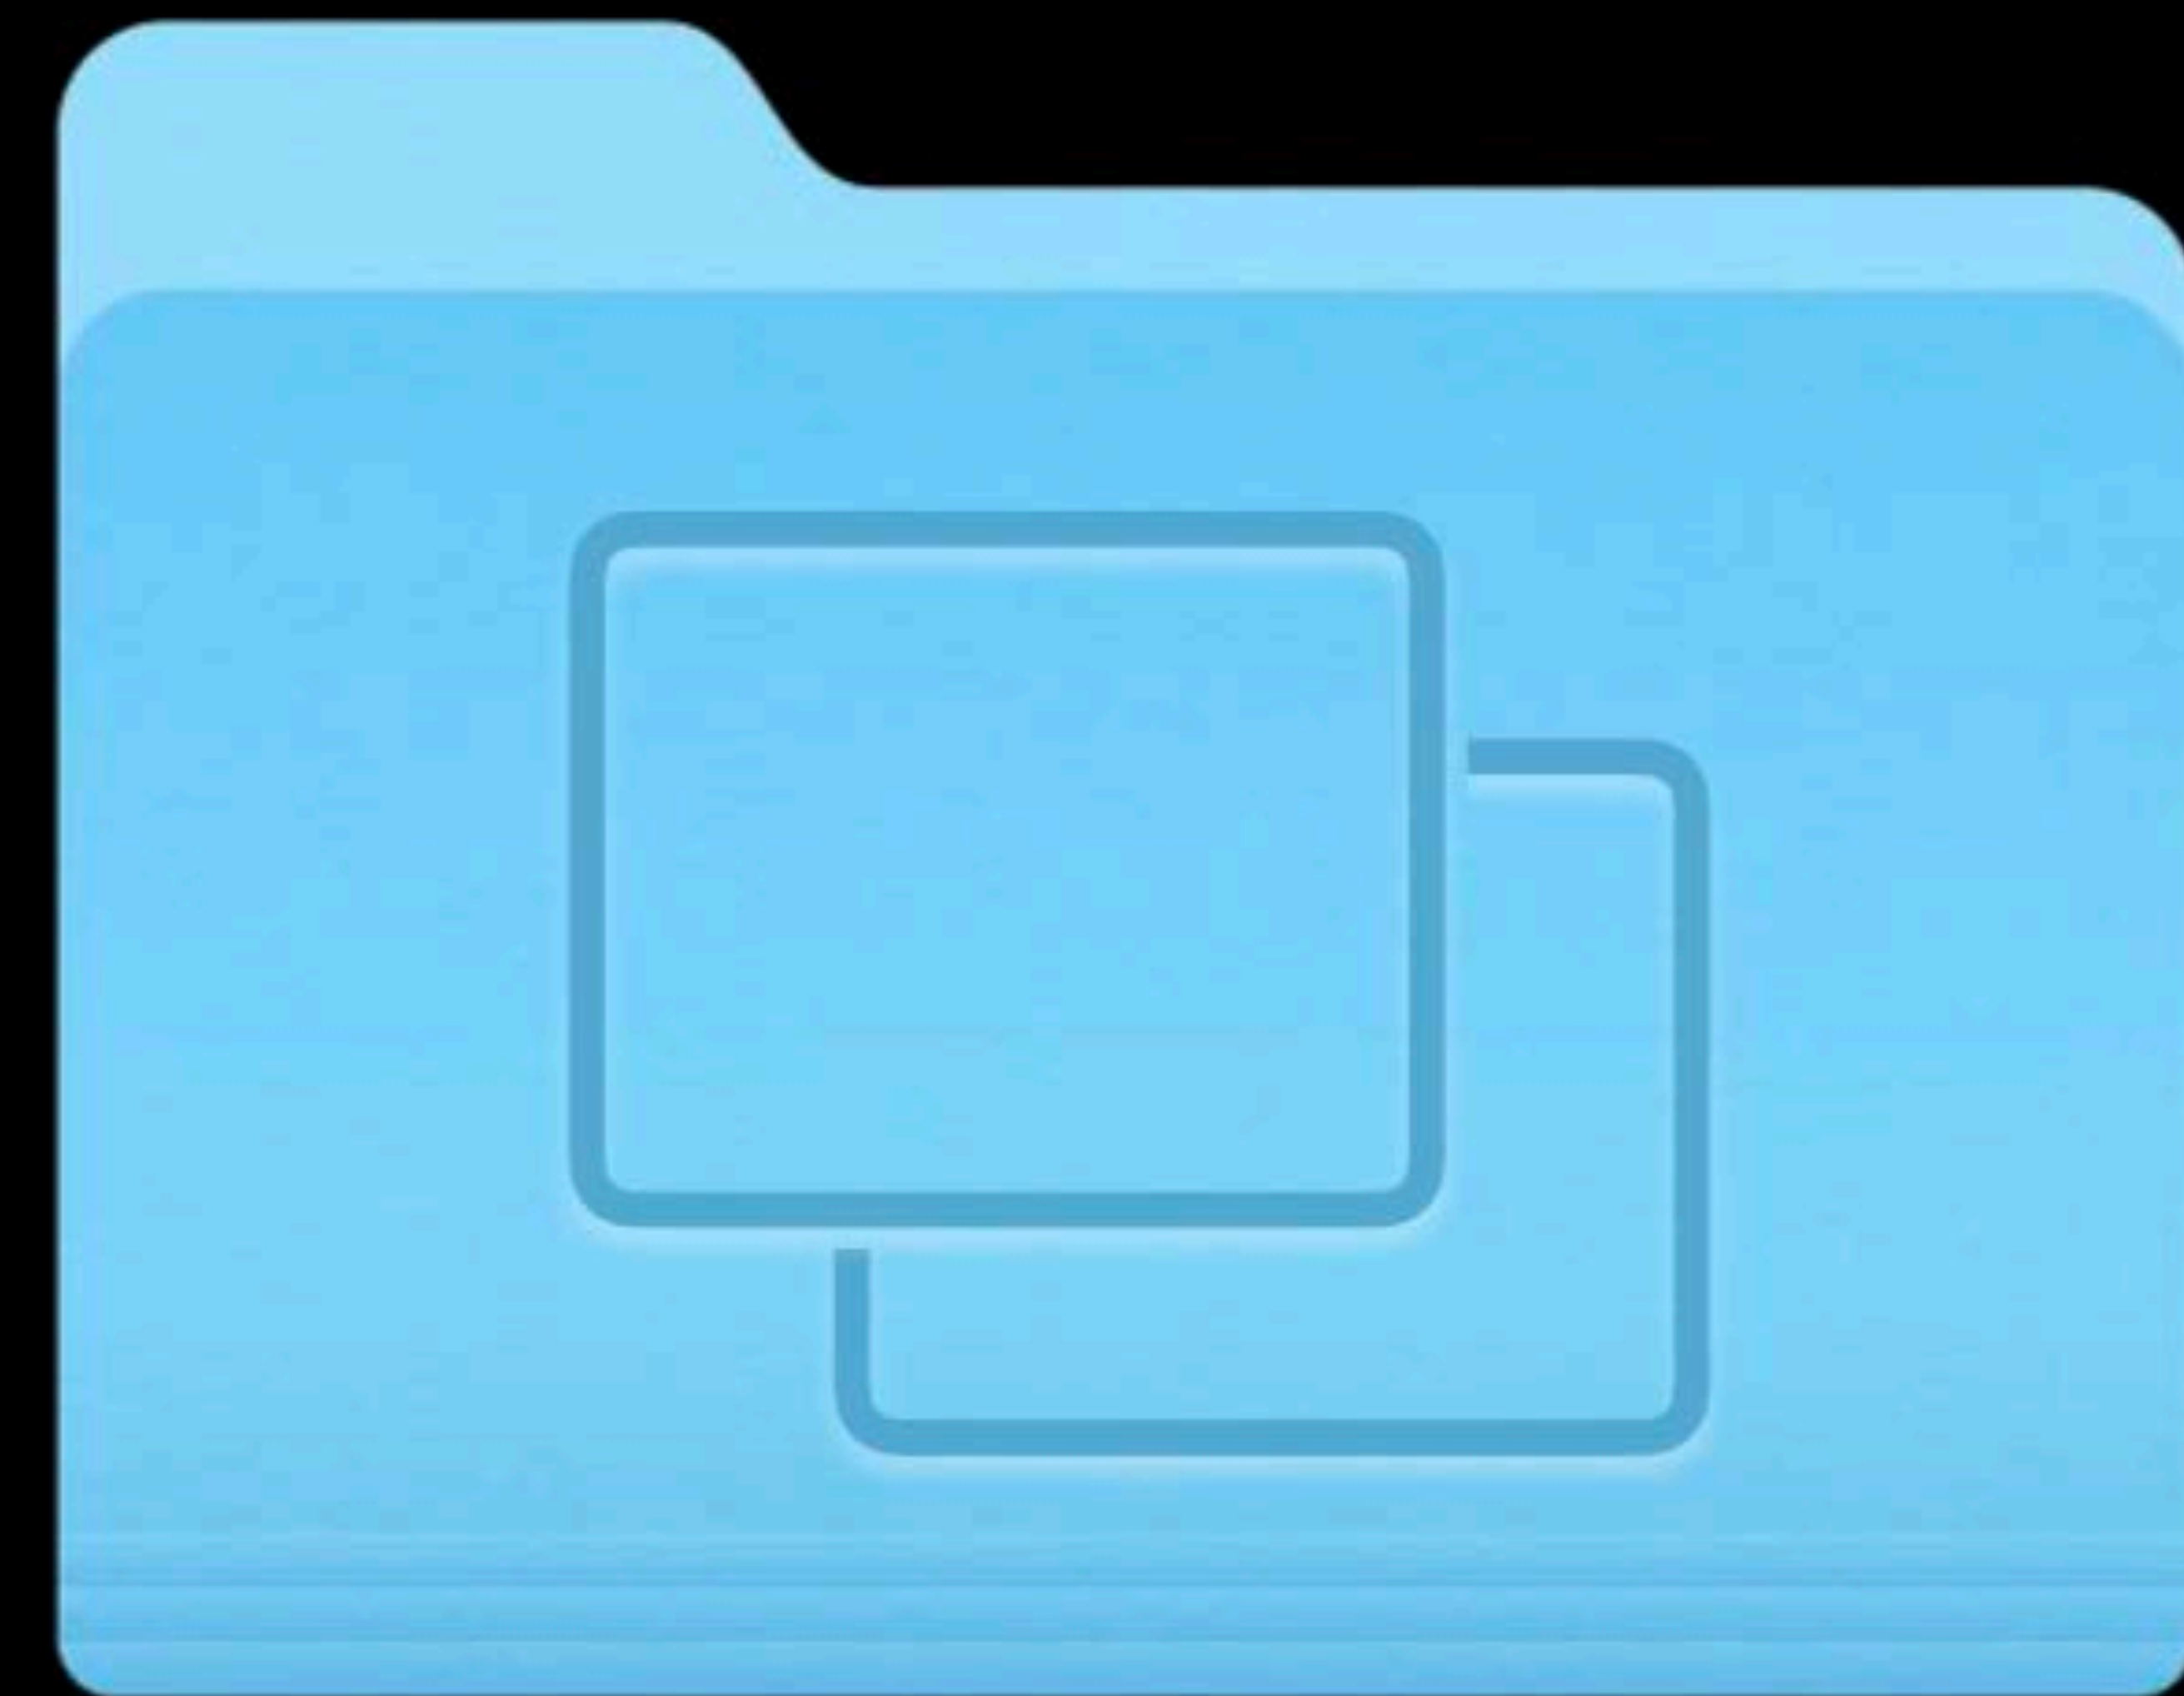
- Image loading
- App launch (up to 10% faster!)



# Asset Catalog

Easy GPU based compression

Lossless by default

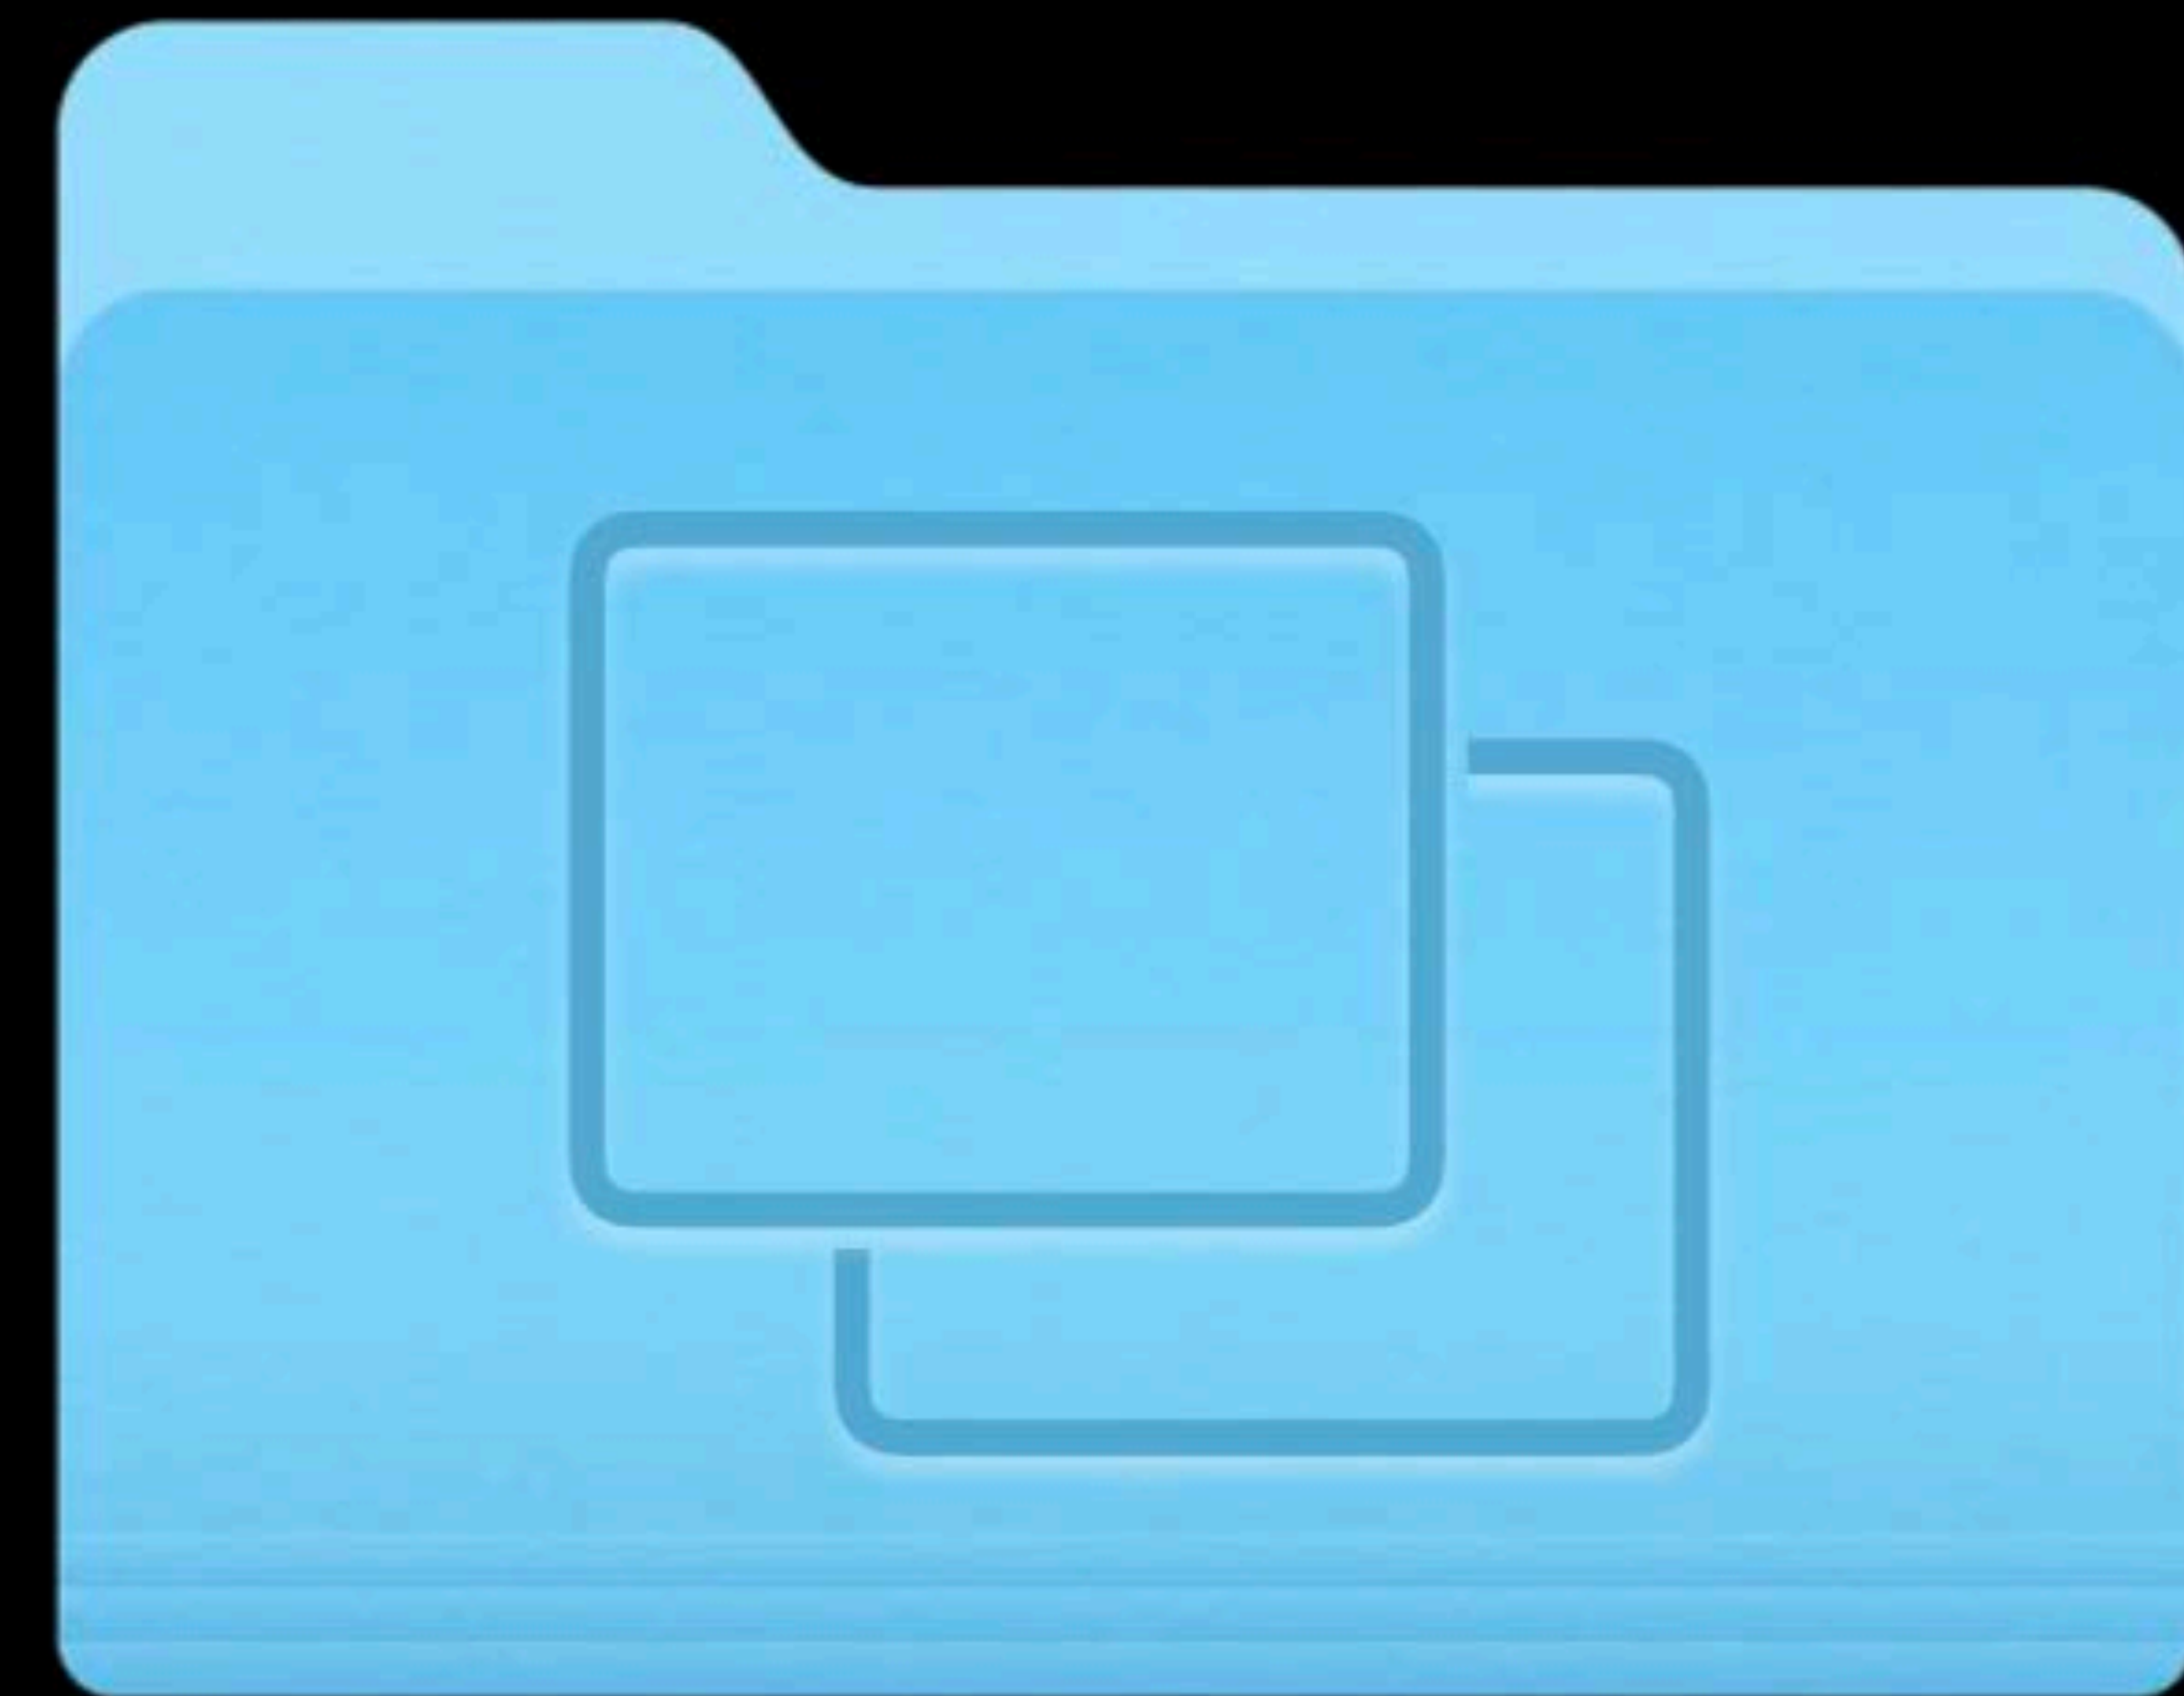


# Asset Catalog

Easy GPU based compression

Lossless by default

Lossy image compression available





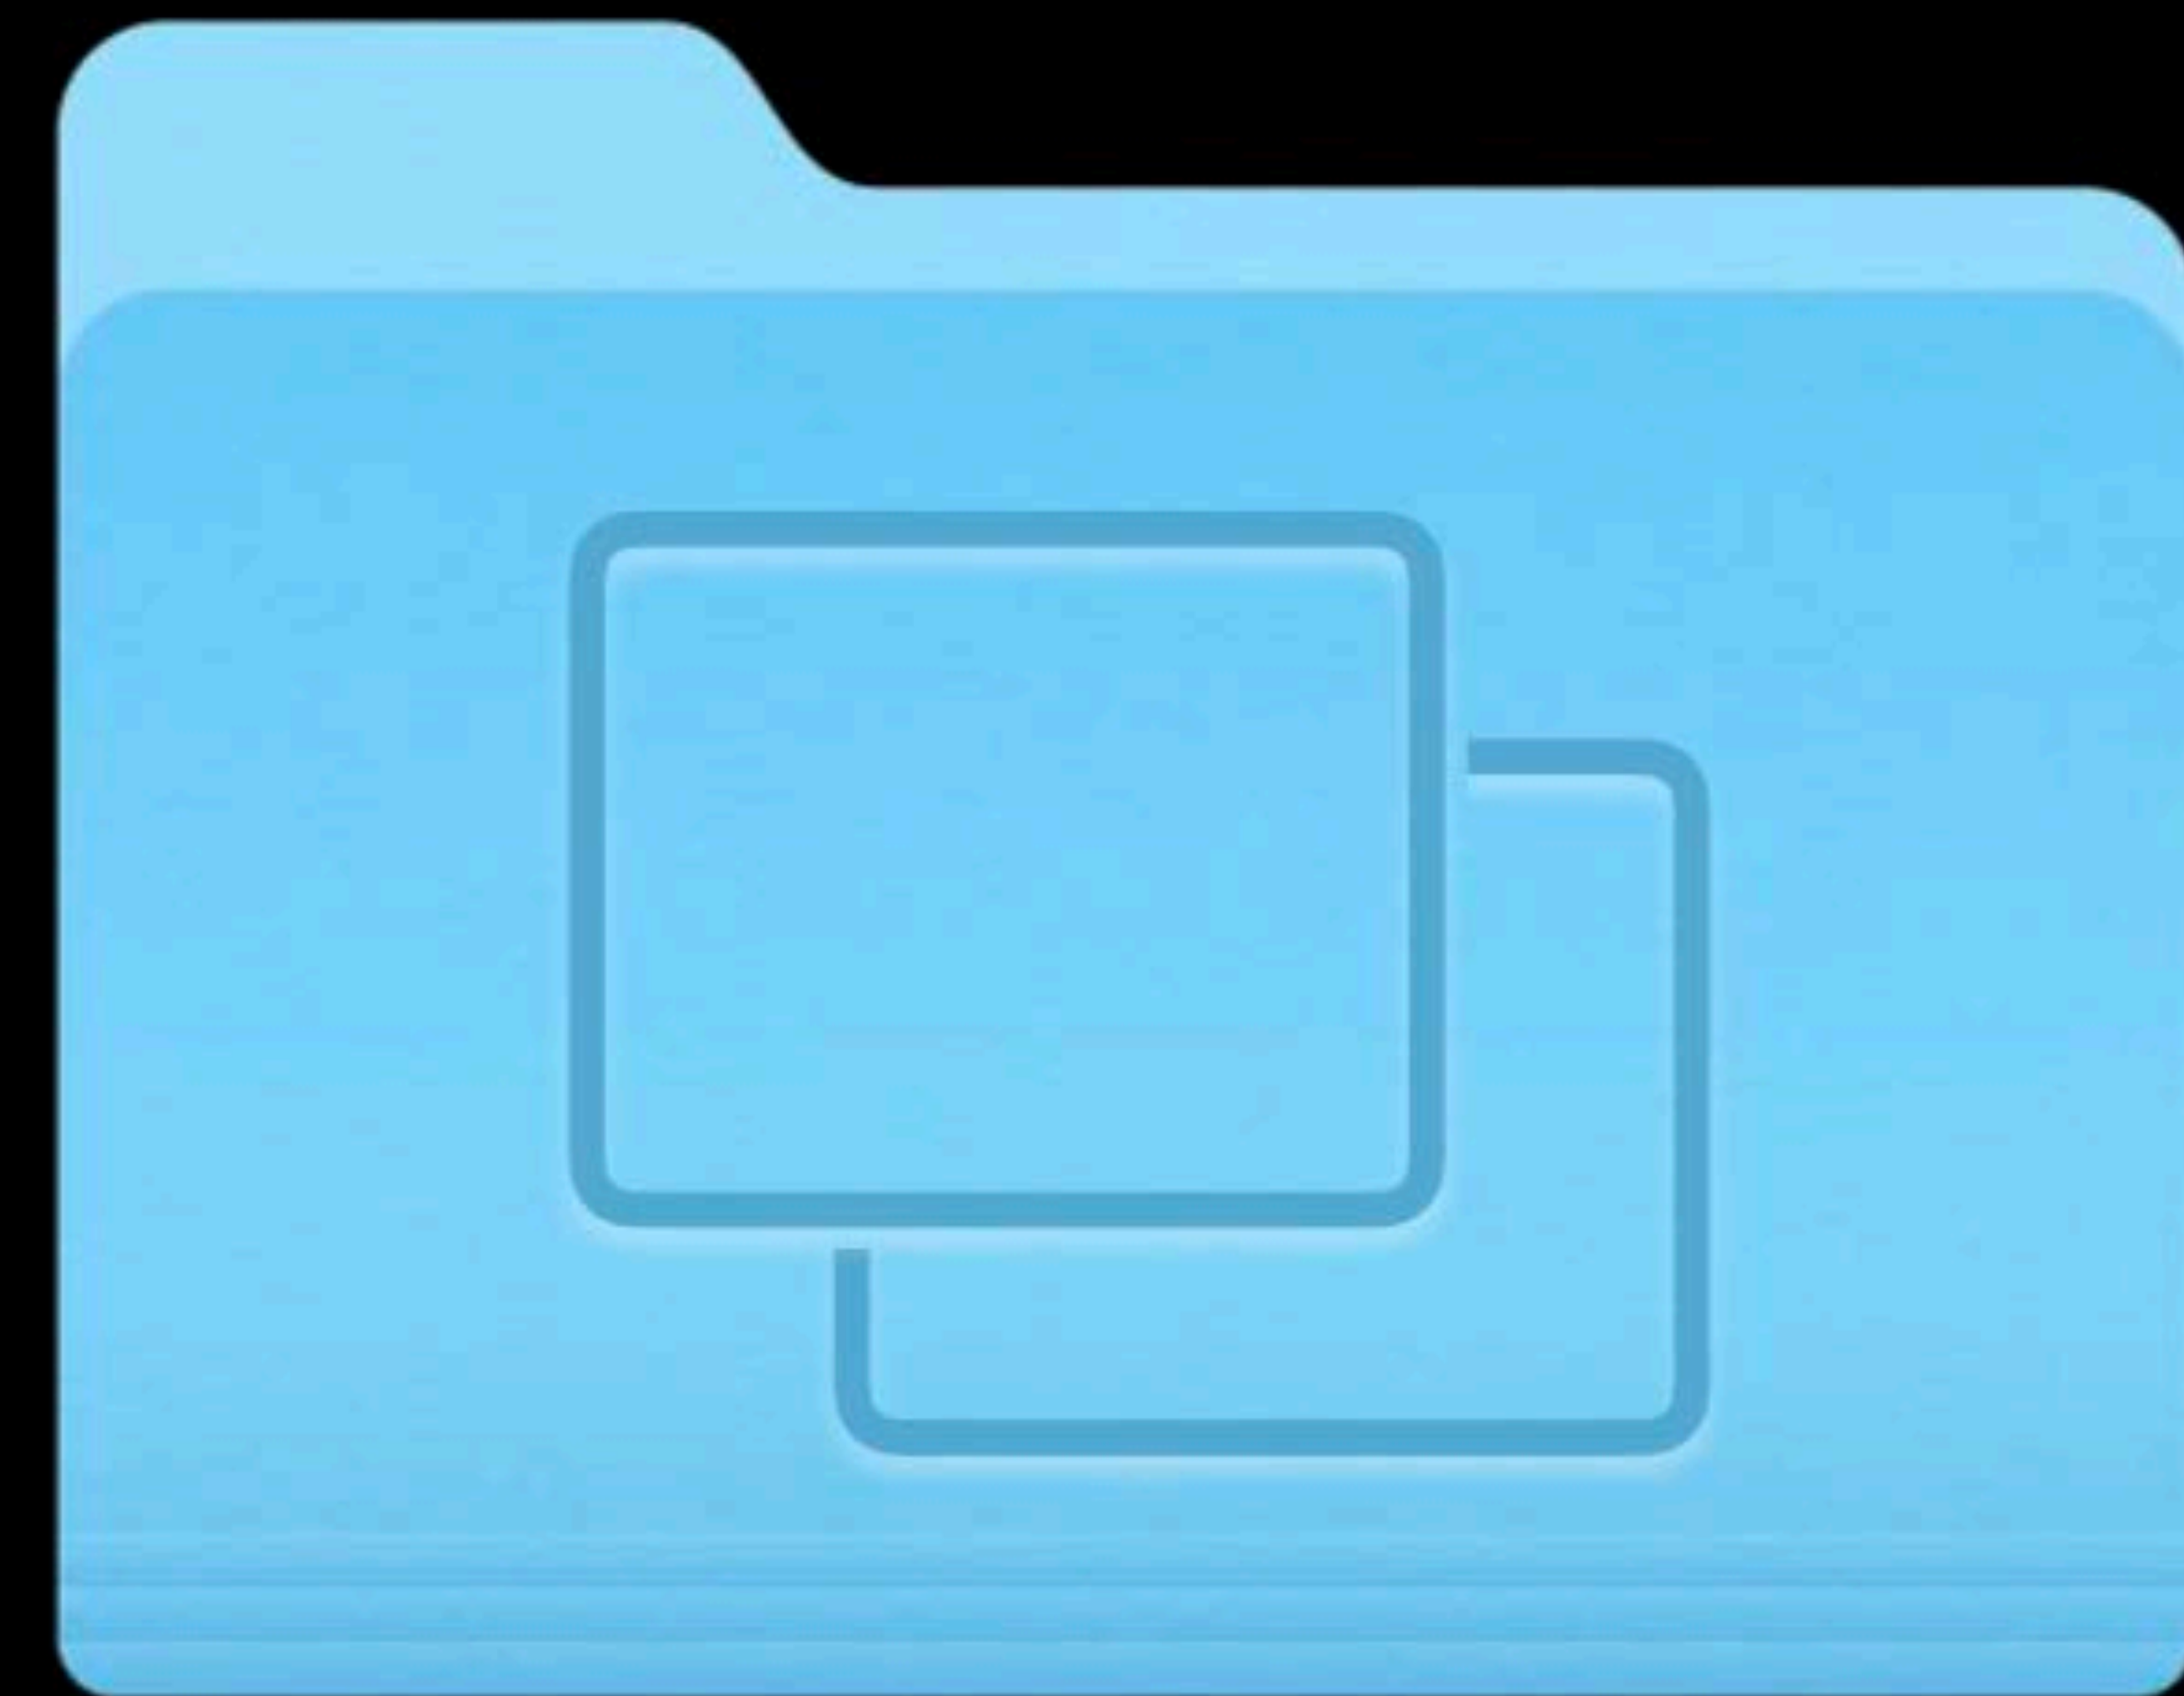
# Asset Catalog

Easy GPU based compression

Lossless by default

Lossy image compression available

- Hardware accelerated decompression



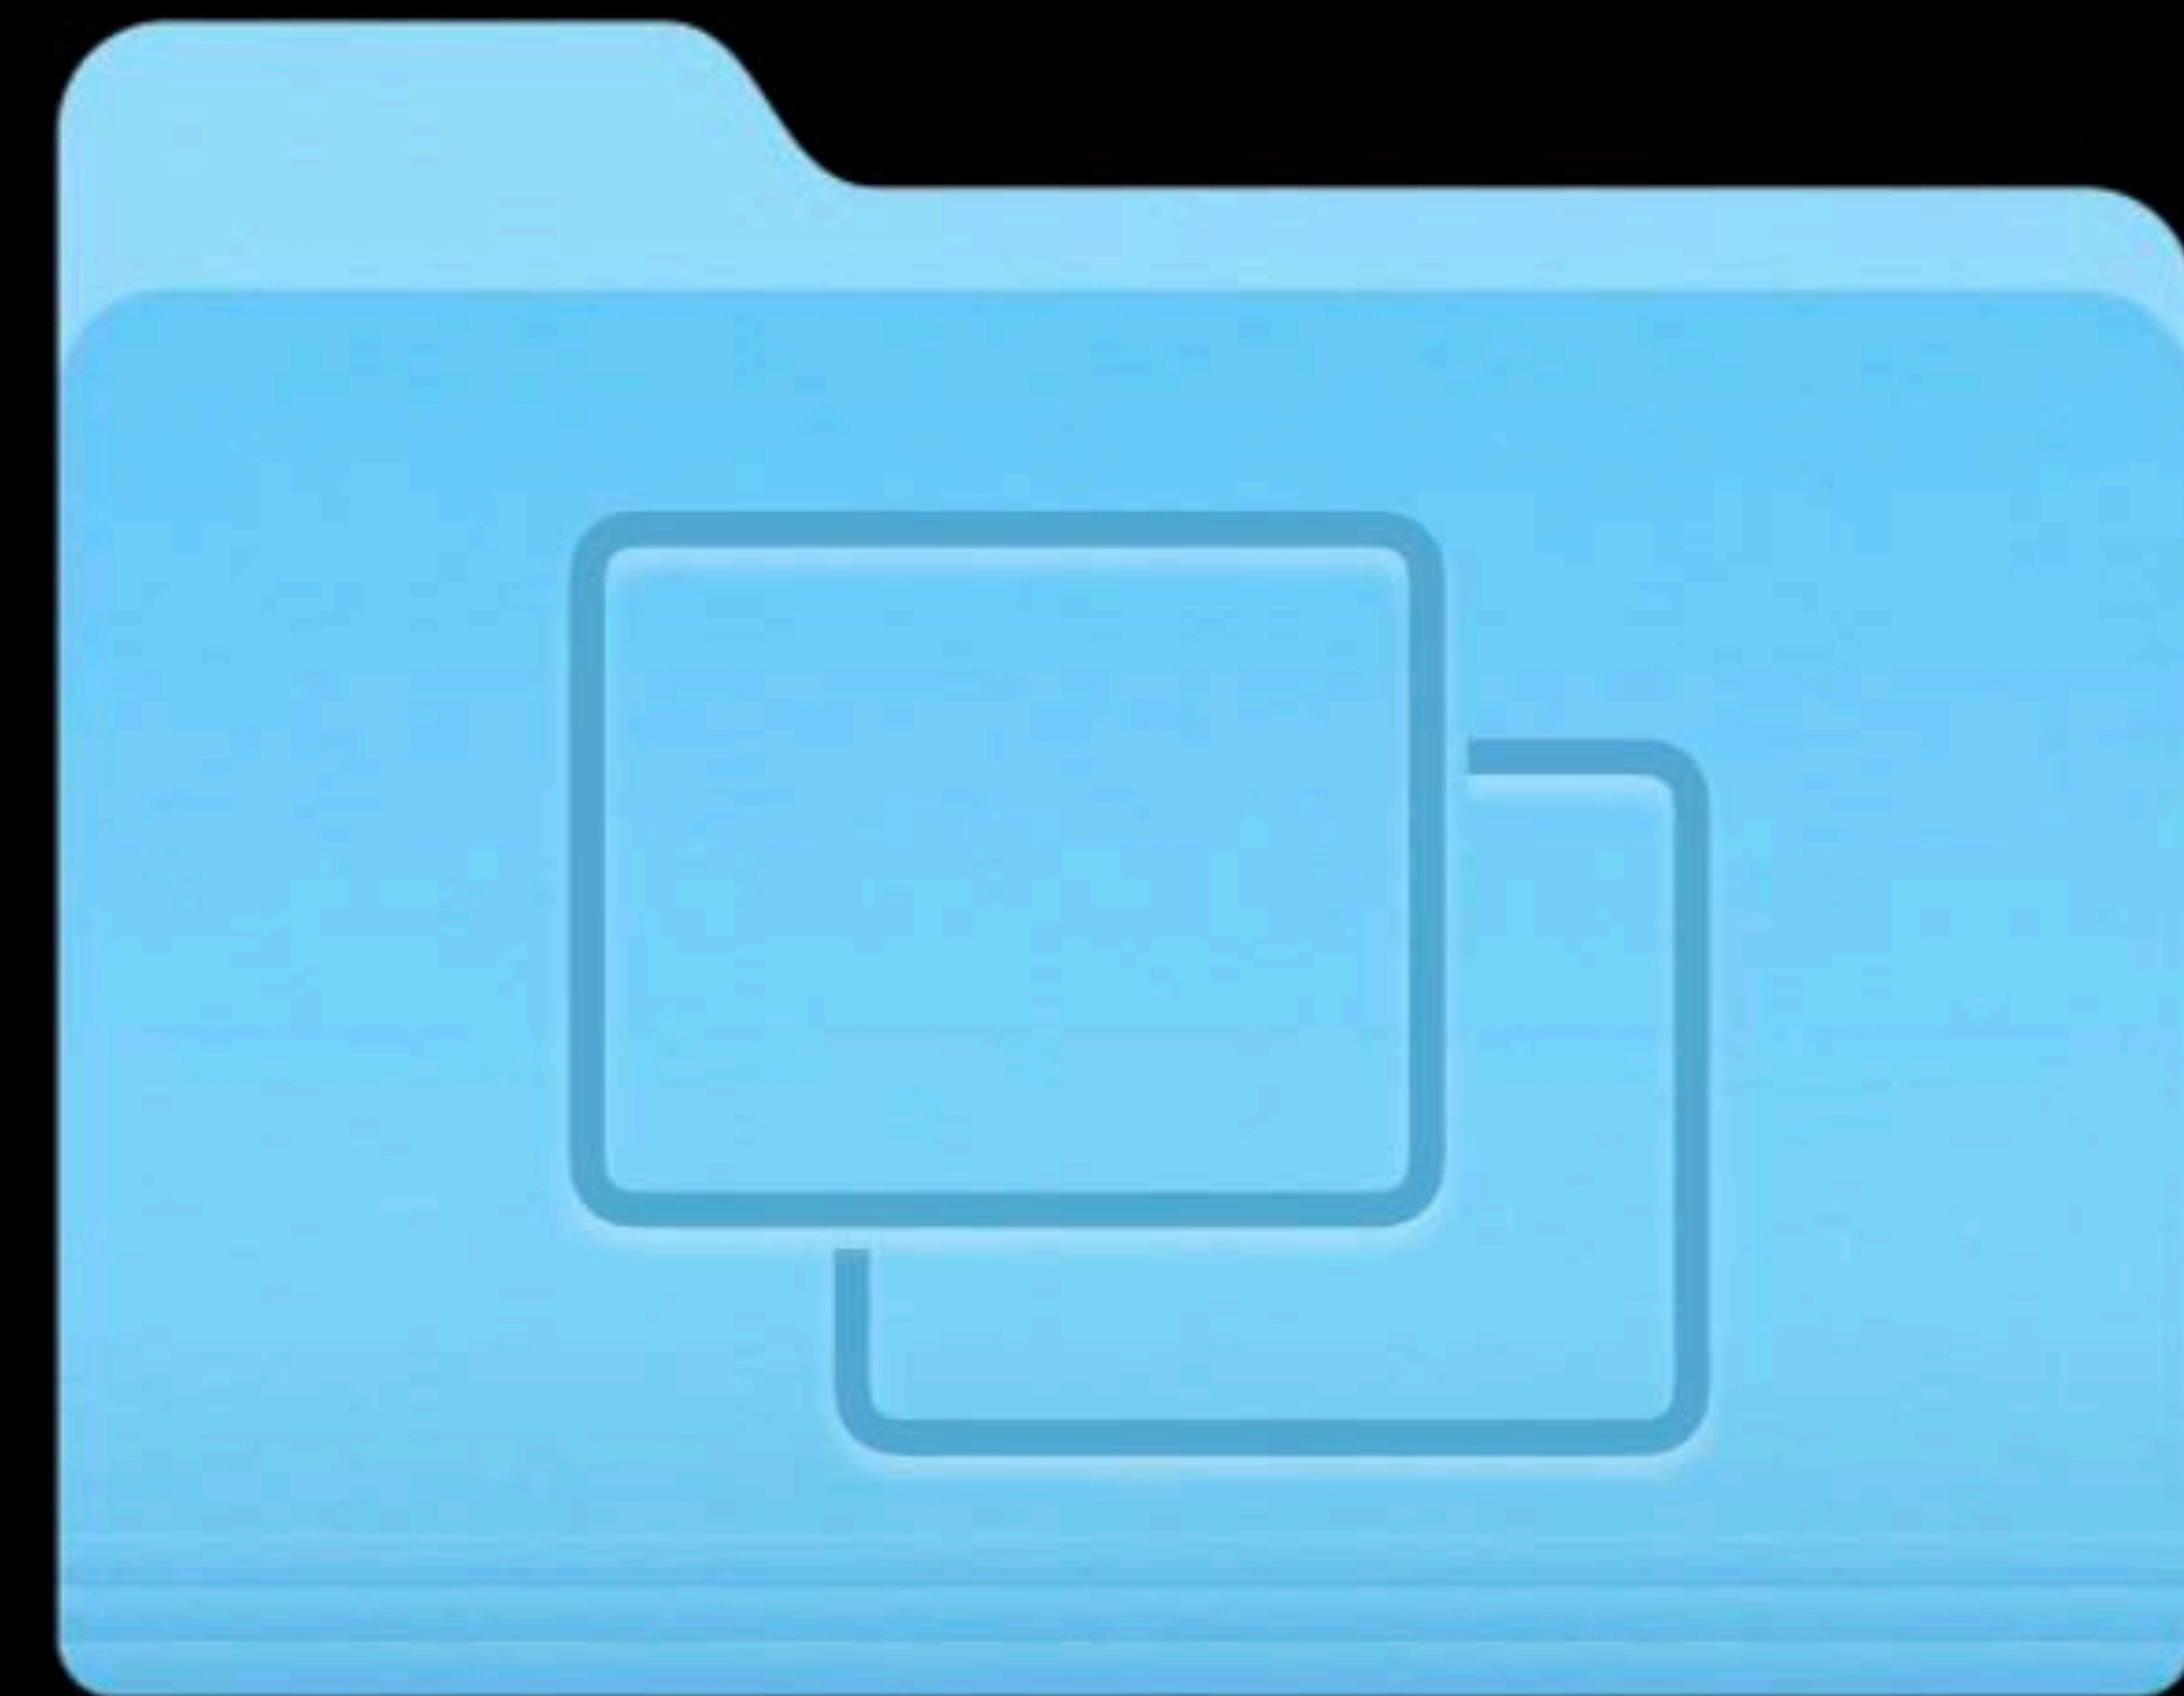
# Asset Catalog

Easy GPU based compression

Lossless by default

Lossy image compression available

- Hardware accelerated decompression
- Lower memory footprint



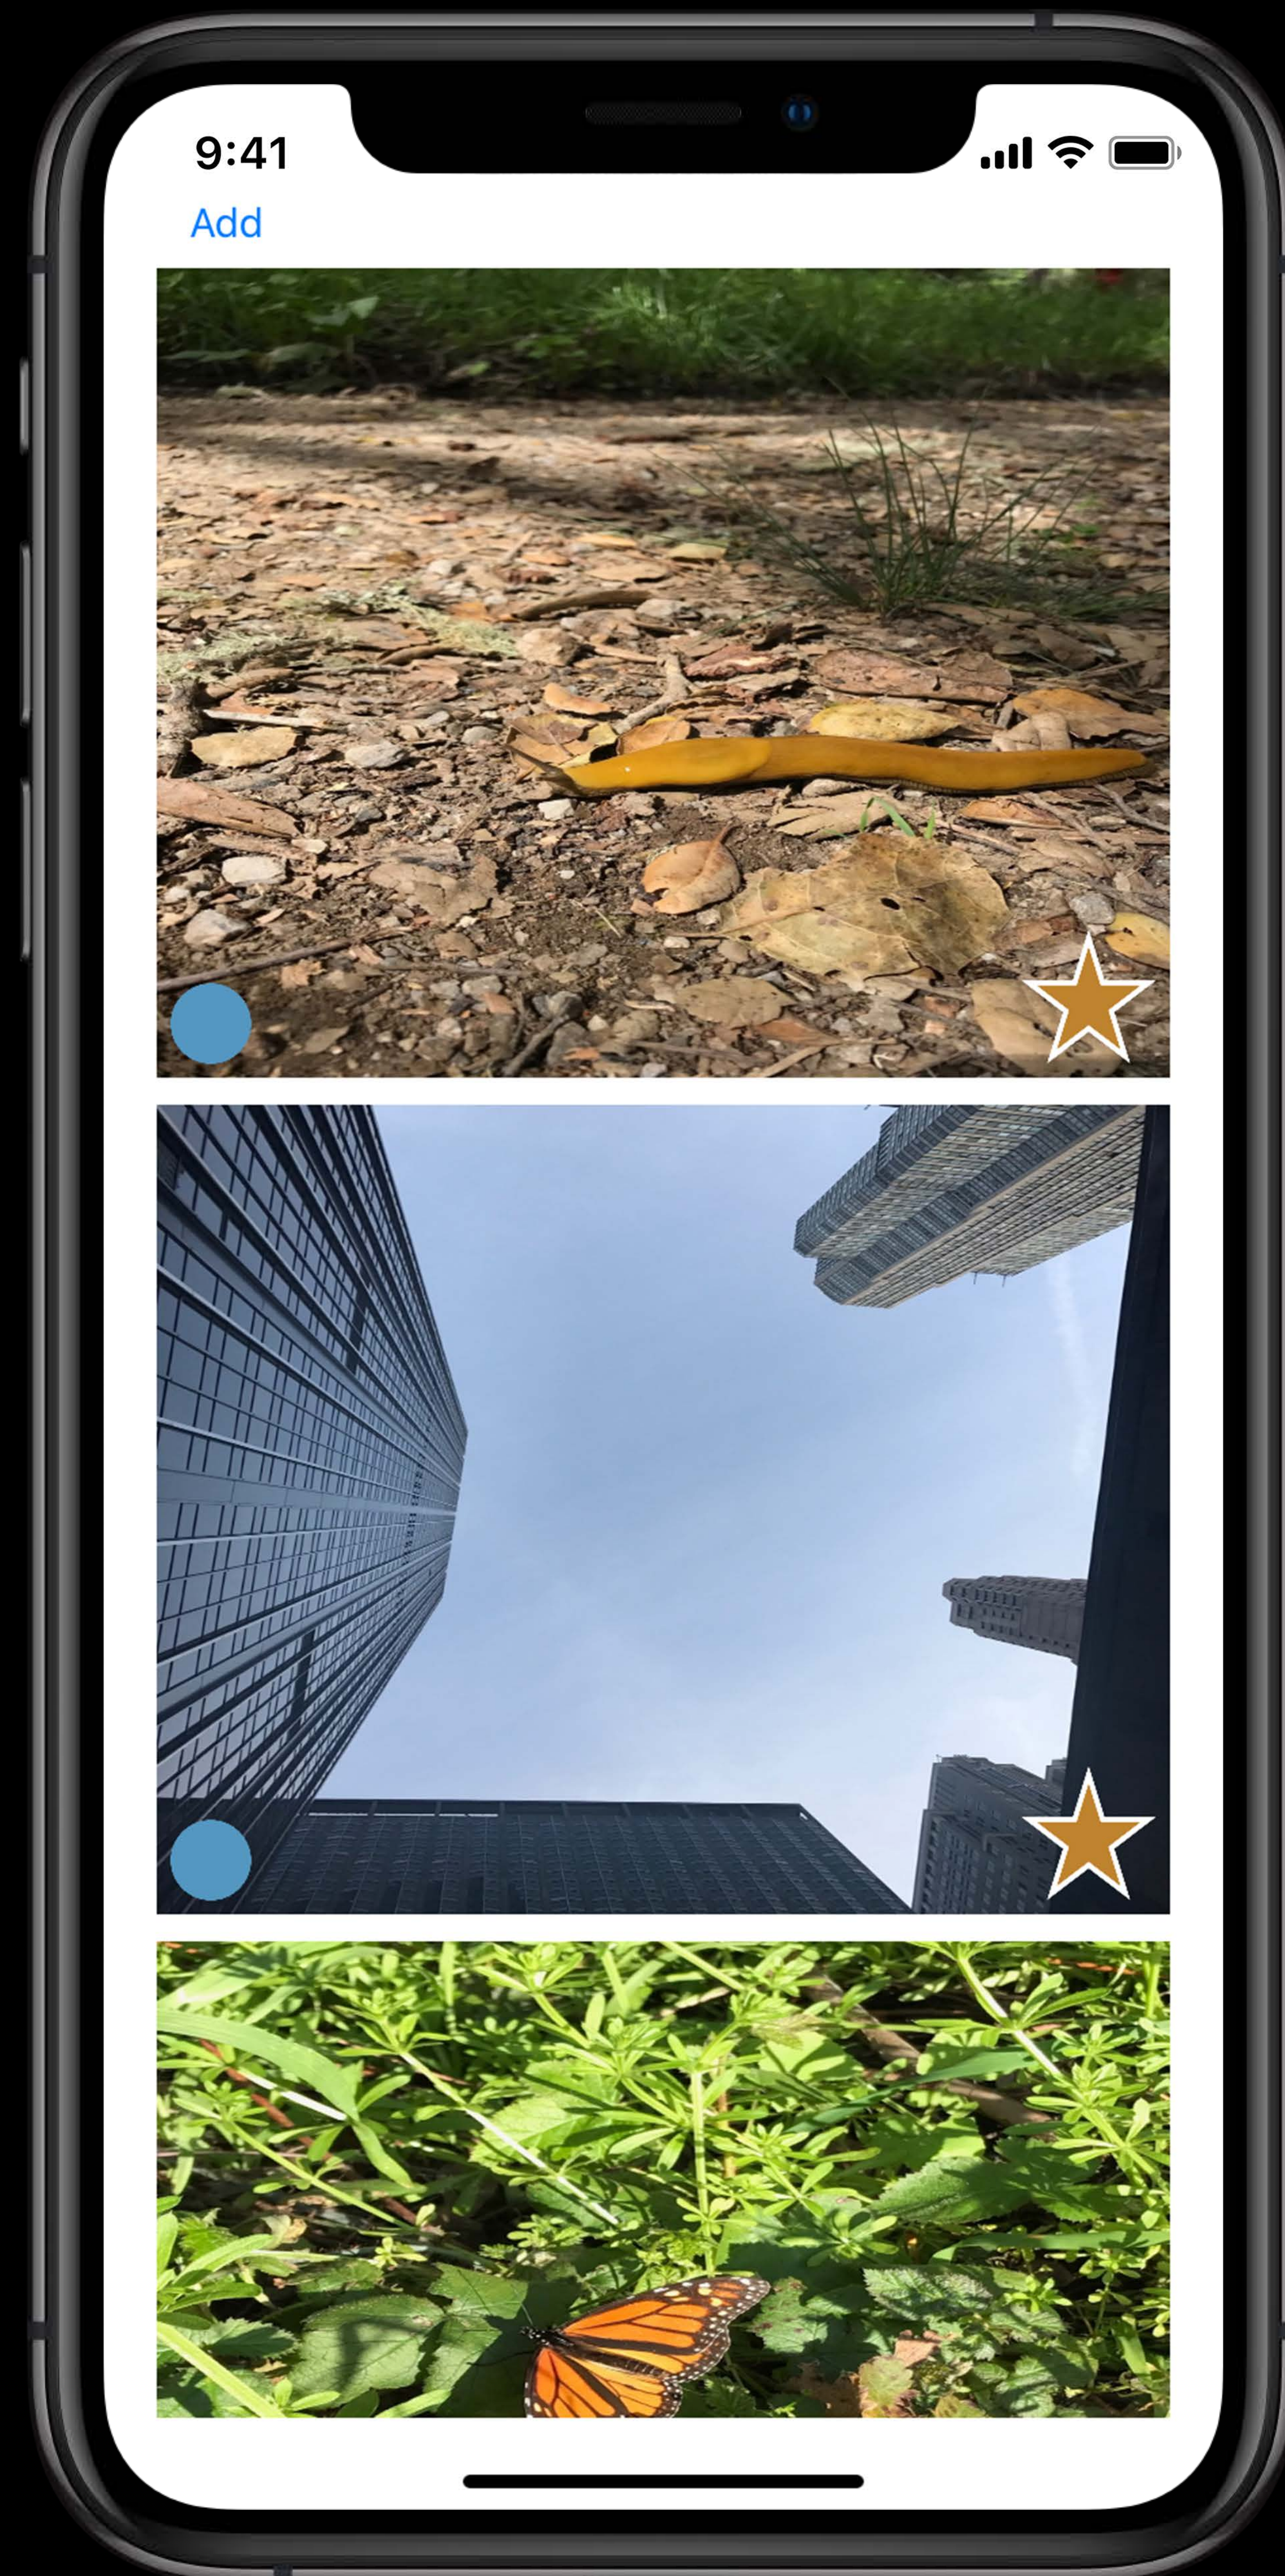


# Image Assets

Growing in size

JPEG Size: 24.6MB

HEIC Size: 17.9MB





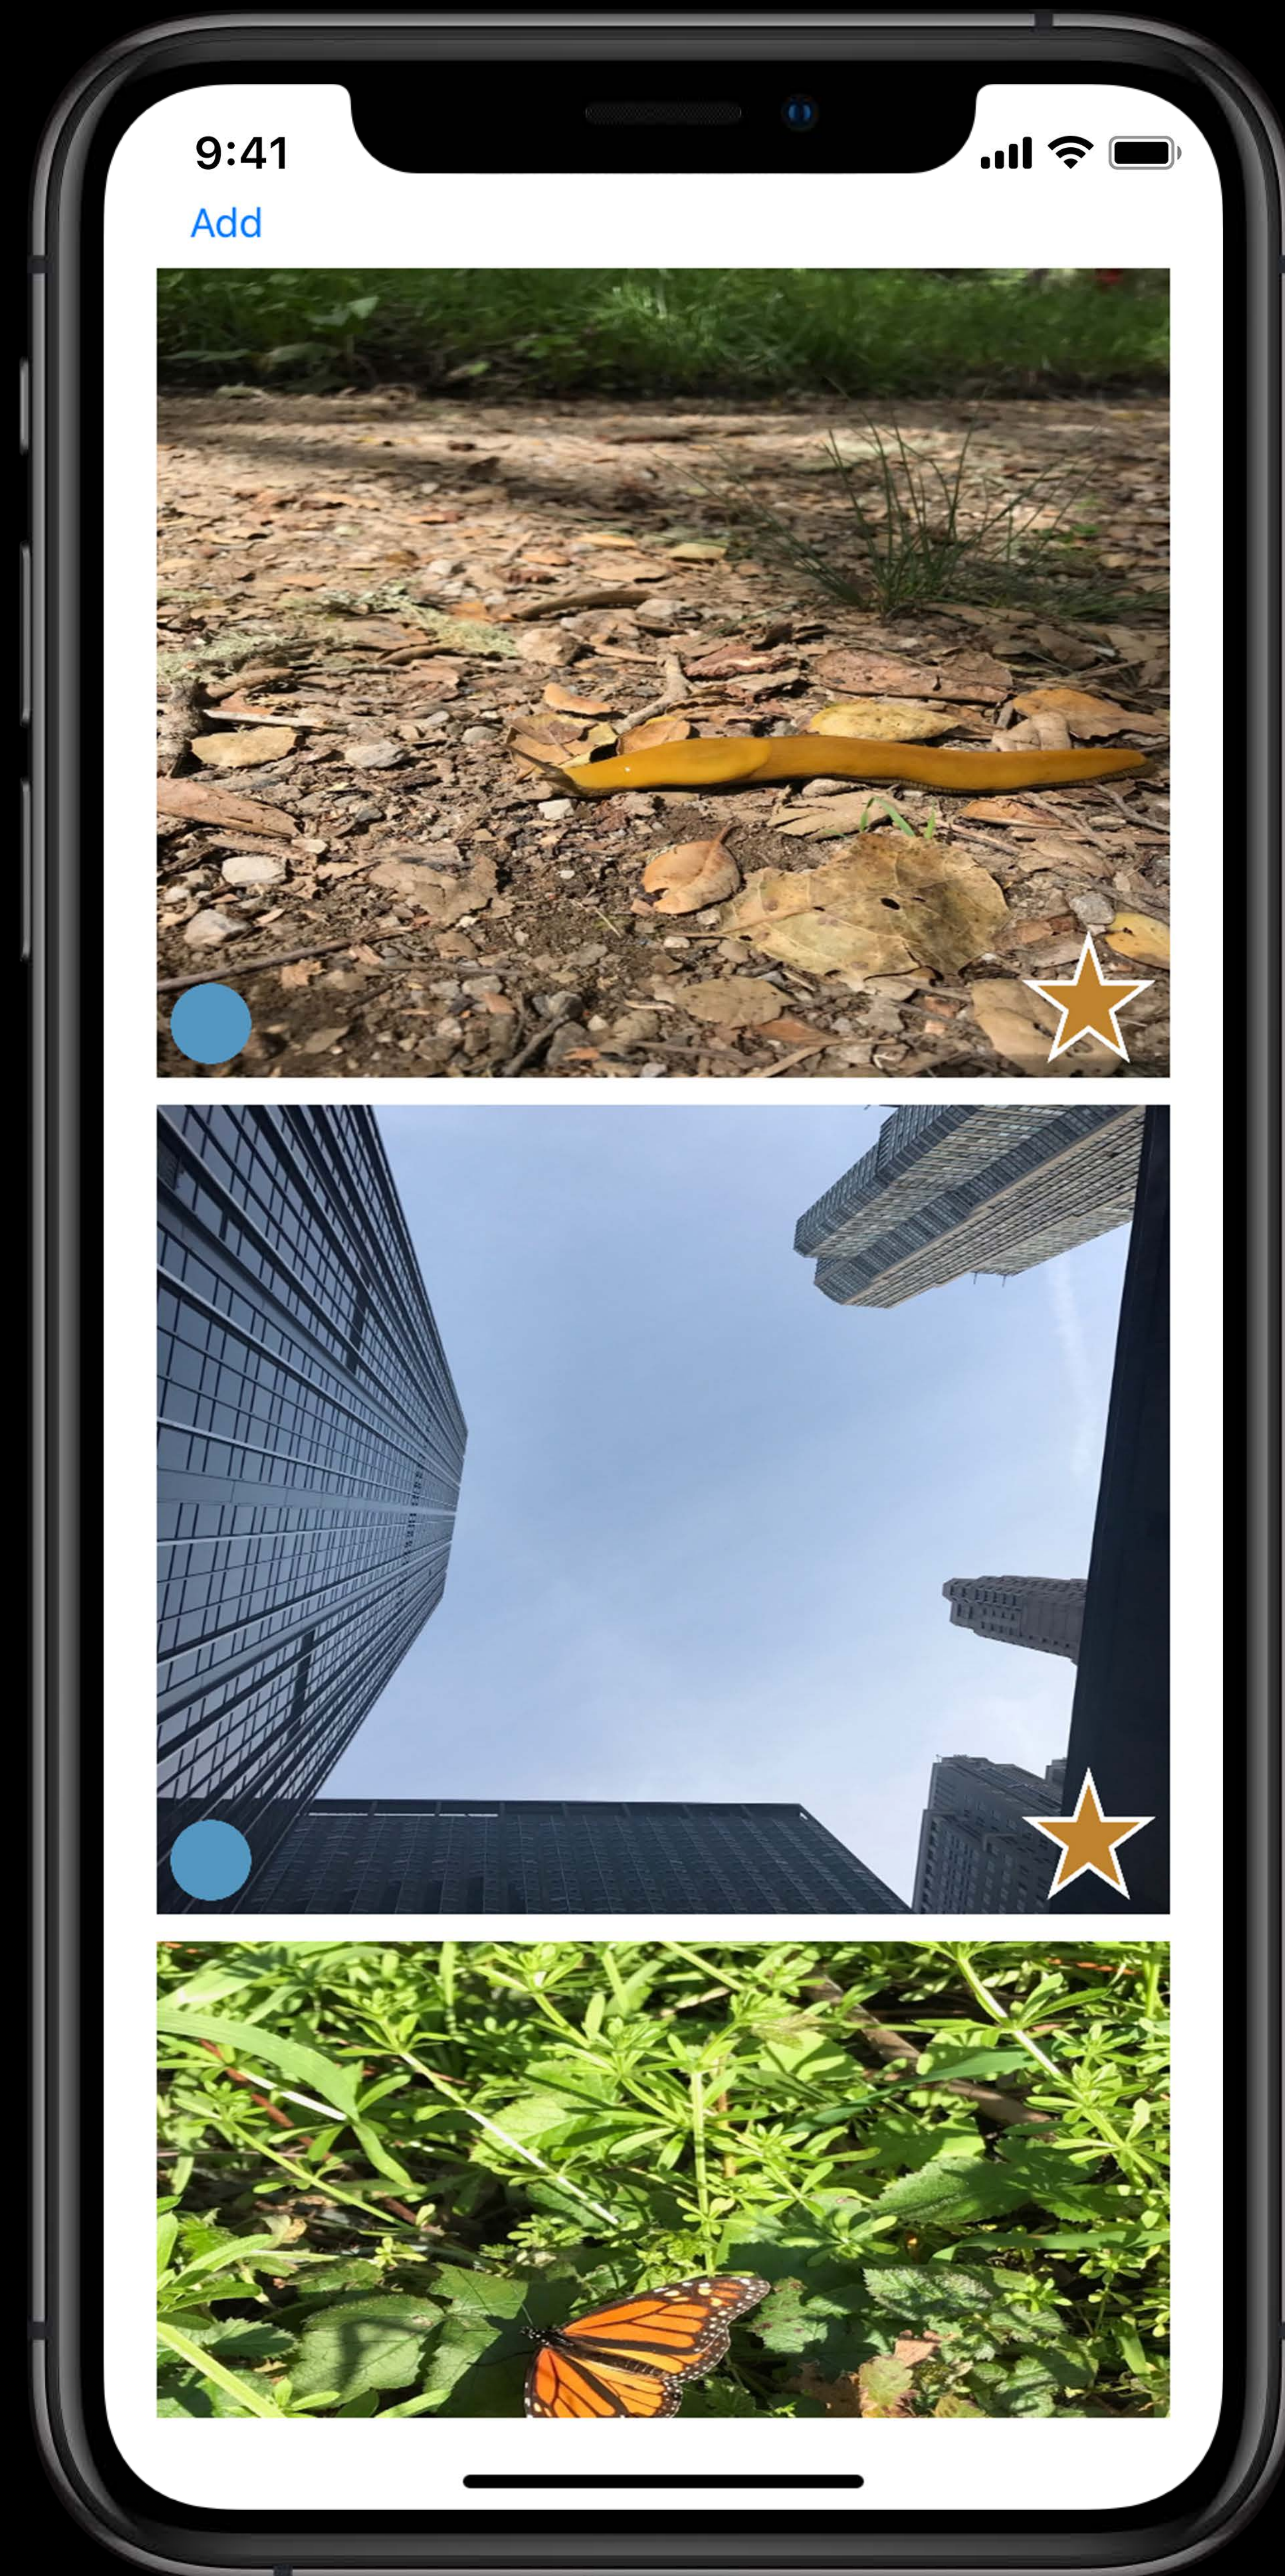
# Image Assets

Growing in size

JPEG Size: 24.6MB

HEIC Size: 17.9MB

AssetCatalog: 14.9MB





Use HEIC and Asset Catalogs  
to reduce footprint

# File System Metadata

# File System Metadata

## Modifying a plist

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN" "http://www.apple.com/DTDs/
PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
  <key>last_app_launch</key>
  <date>2019-06-07T07:26:49Z</date>
</dict>
</plist>
```

# File System Metadata

Modifying a plist

On launch our demo app reads and updates this plist

# File System Metadata

Modifying a plist

On launch our demo app reads and updates this plist

- One read

# File System Metadata

Modifying a plist

On launch our demo app reads and updates this plist

- One read
- Three writes

# File System Metadata

Modifying a plist

On launch our demo app reads and updates this plist

- One read
- Three writes
- `fsync`



# File System Metadata

Why three writes instead of one?

File system metadata writes happen as a result of

- Creating a file
- Deleting a file
- Renaming a file
- Updating a file
- And other file system operations

# File System Metadata

What is file system metadata?

Name



# File System Metadata

What is file system metadata?

Name

Size

# File System Metadata

What is file system metadata?

Name

Size

Location



# File System Metadata

What is file system metadata?

Name

Size

Location

Dates



# File System Metadata

What is file system metadata?

Name

Size

Location

Dates

And more

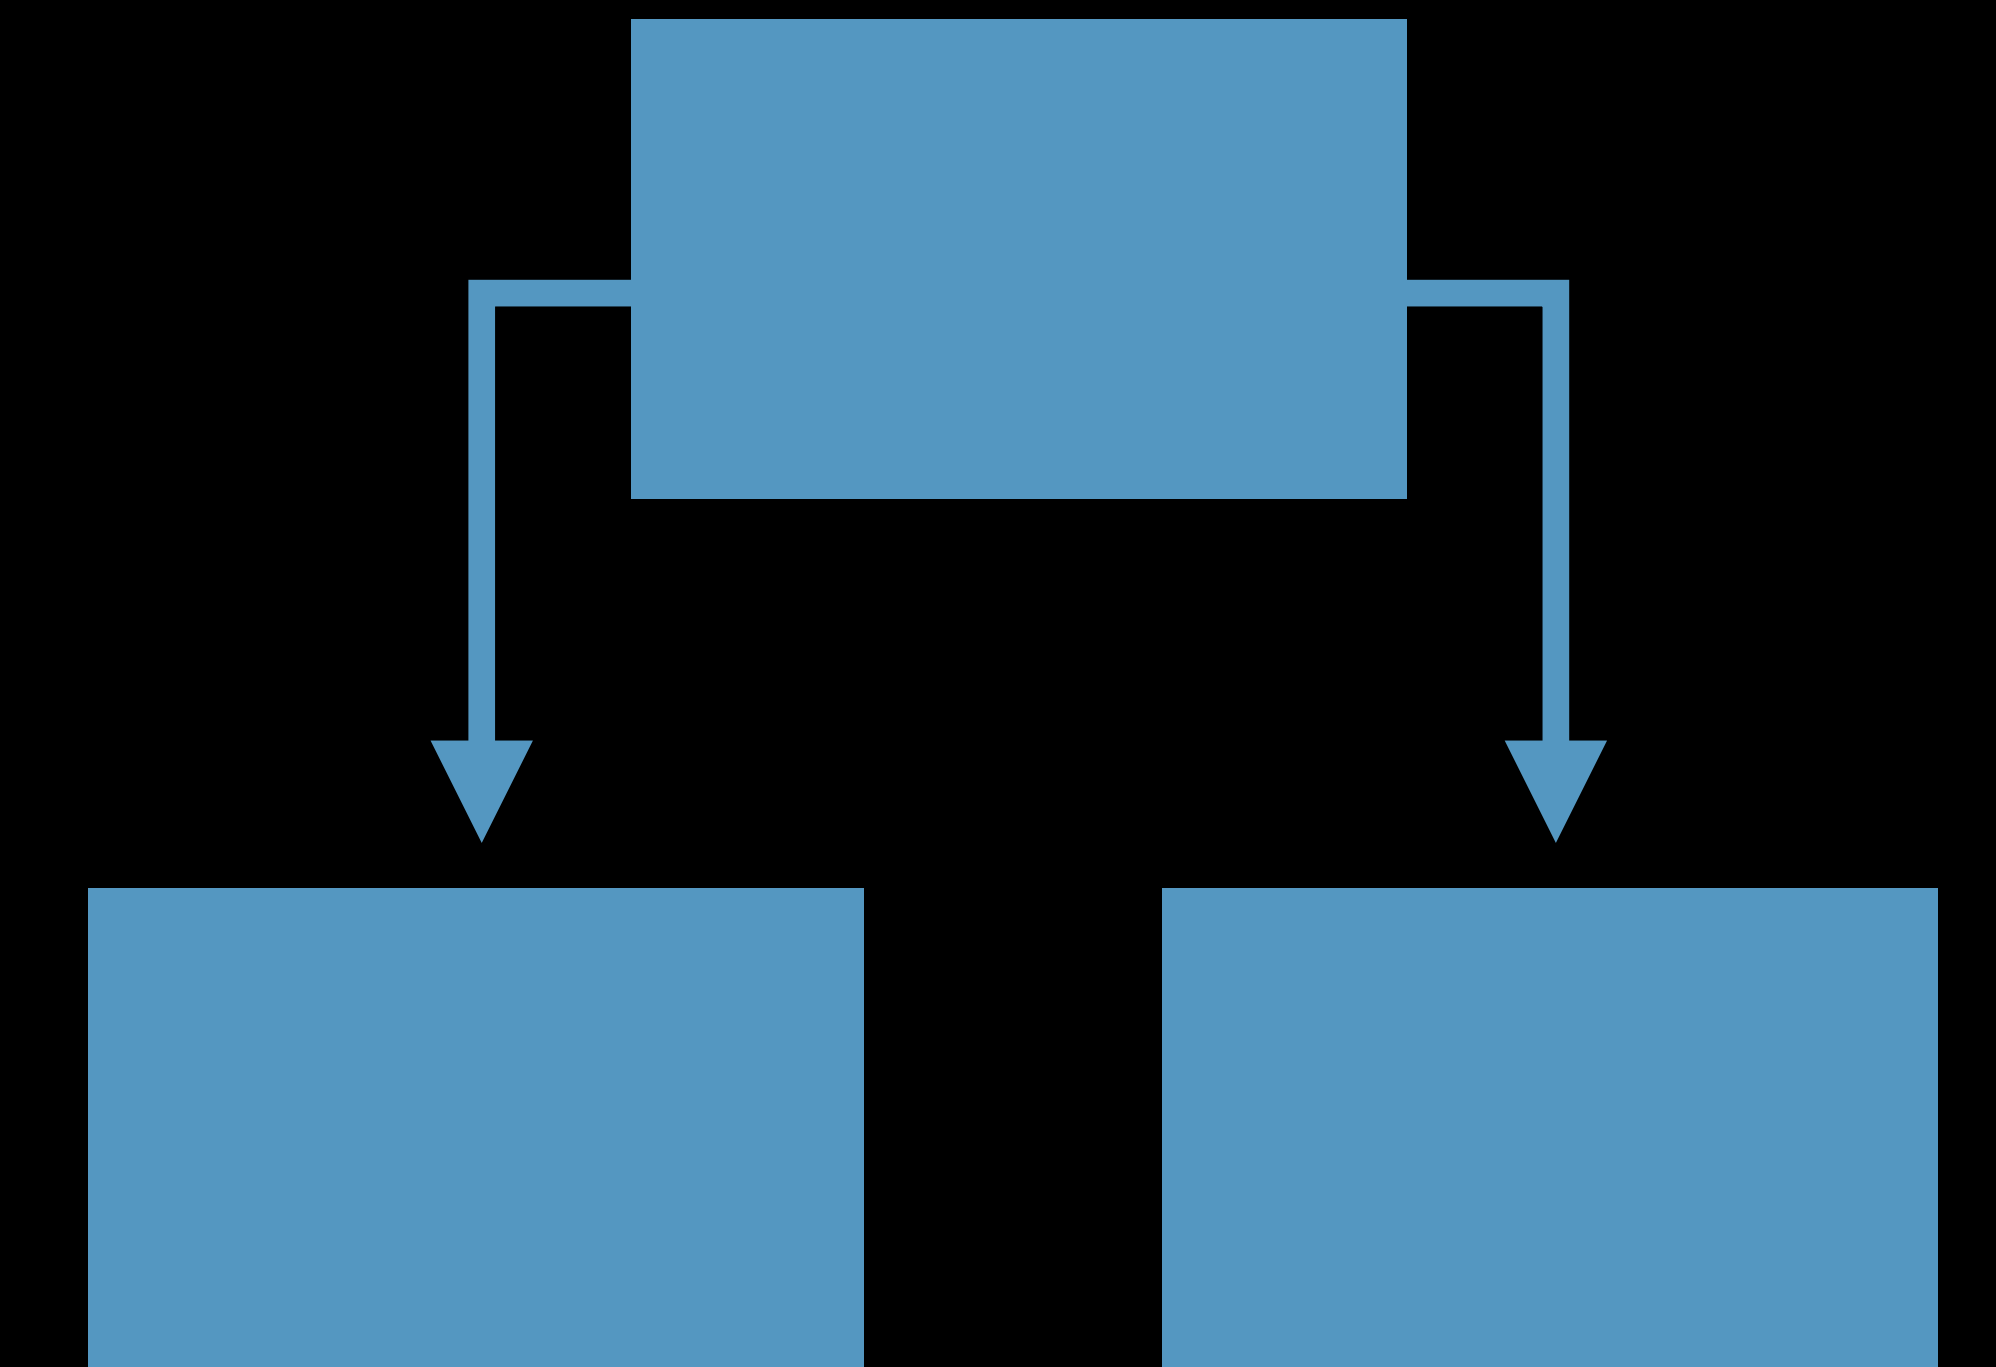
# File System Metadata

Writing a 240 byte NSDictionary to a file

Object Map



File System Tree





# File System Metadata

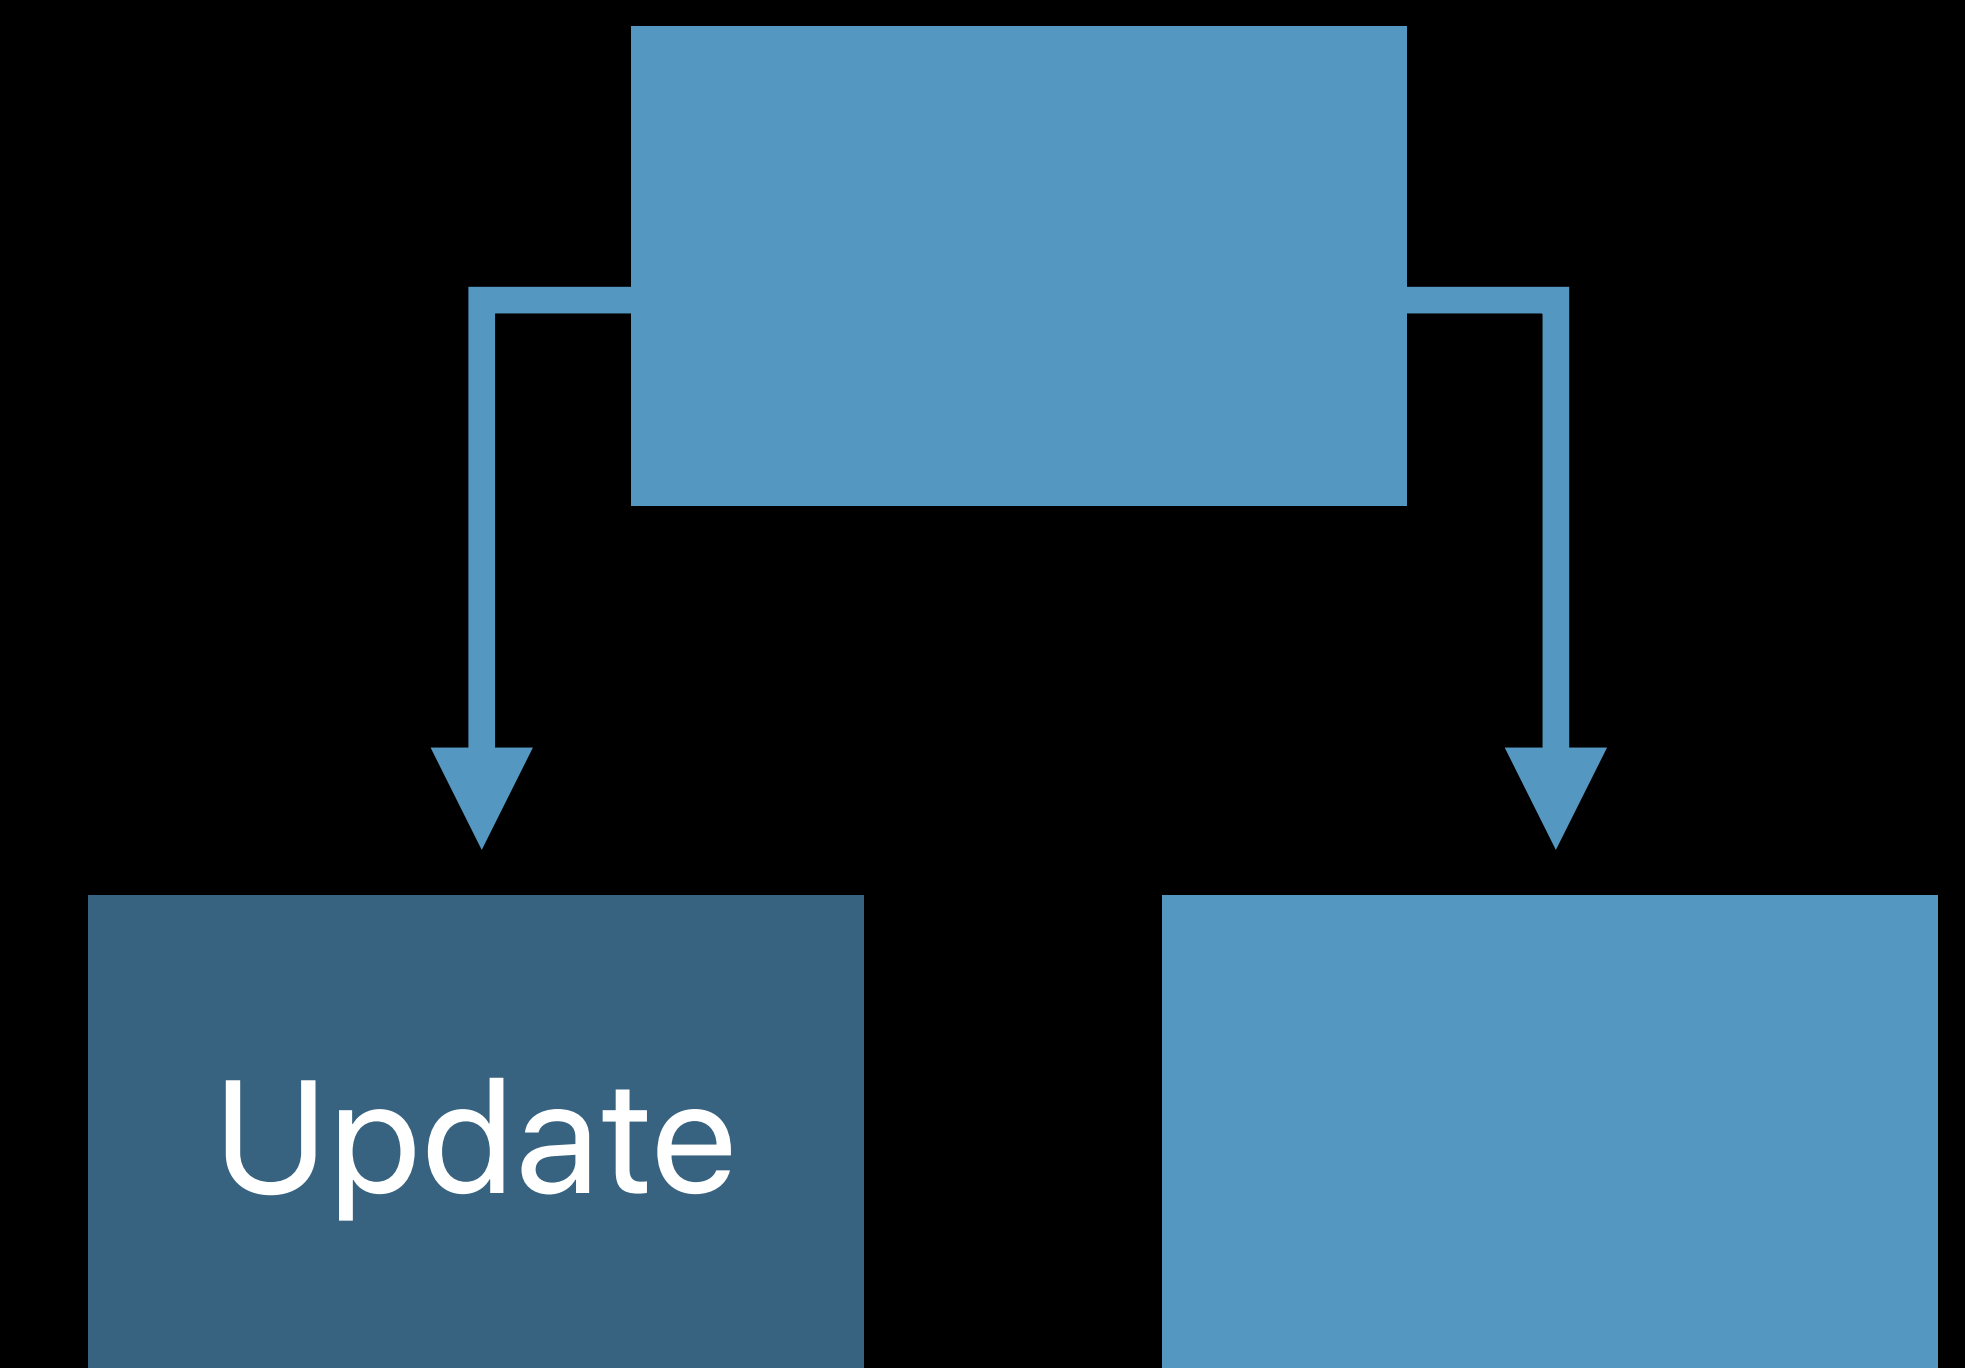
Writing a 240 byte NSDictionary to a file

Update file system node

Object Map



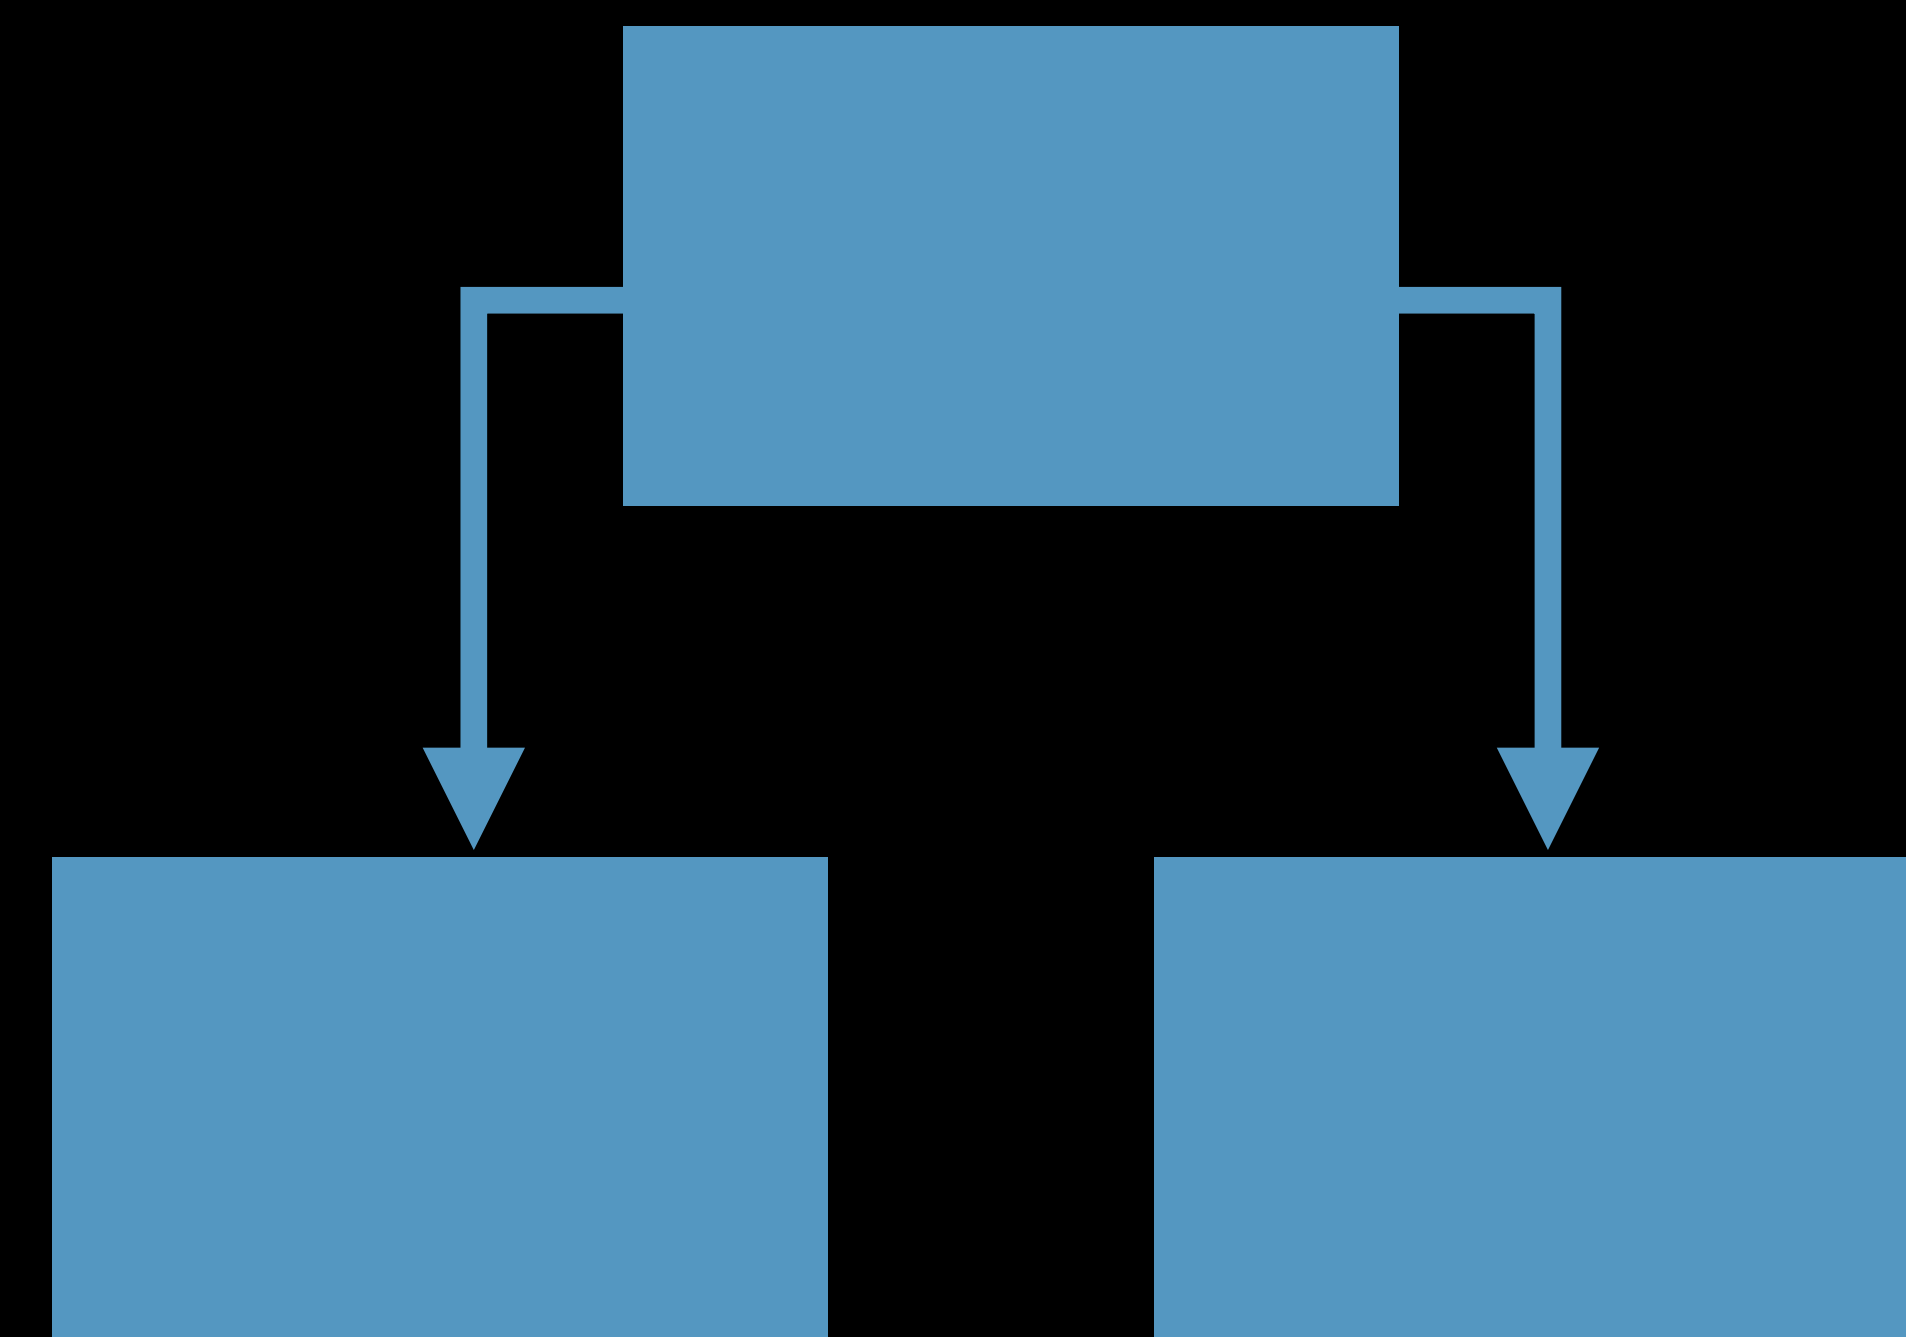
File System Tree



# File System Metadata

Writing a 240 byte NSDictionary to a file

## File System Tree

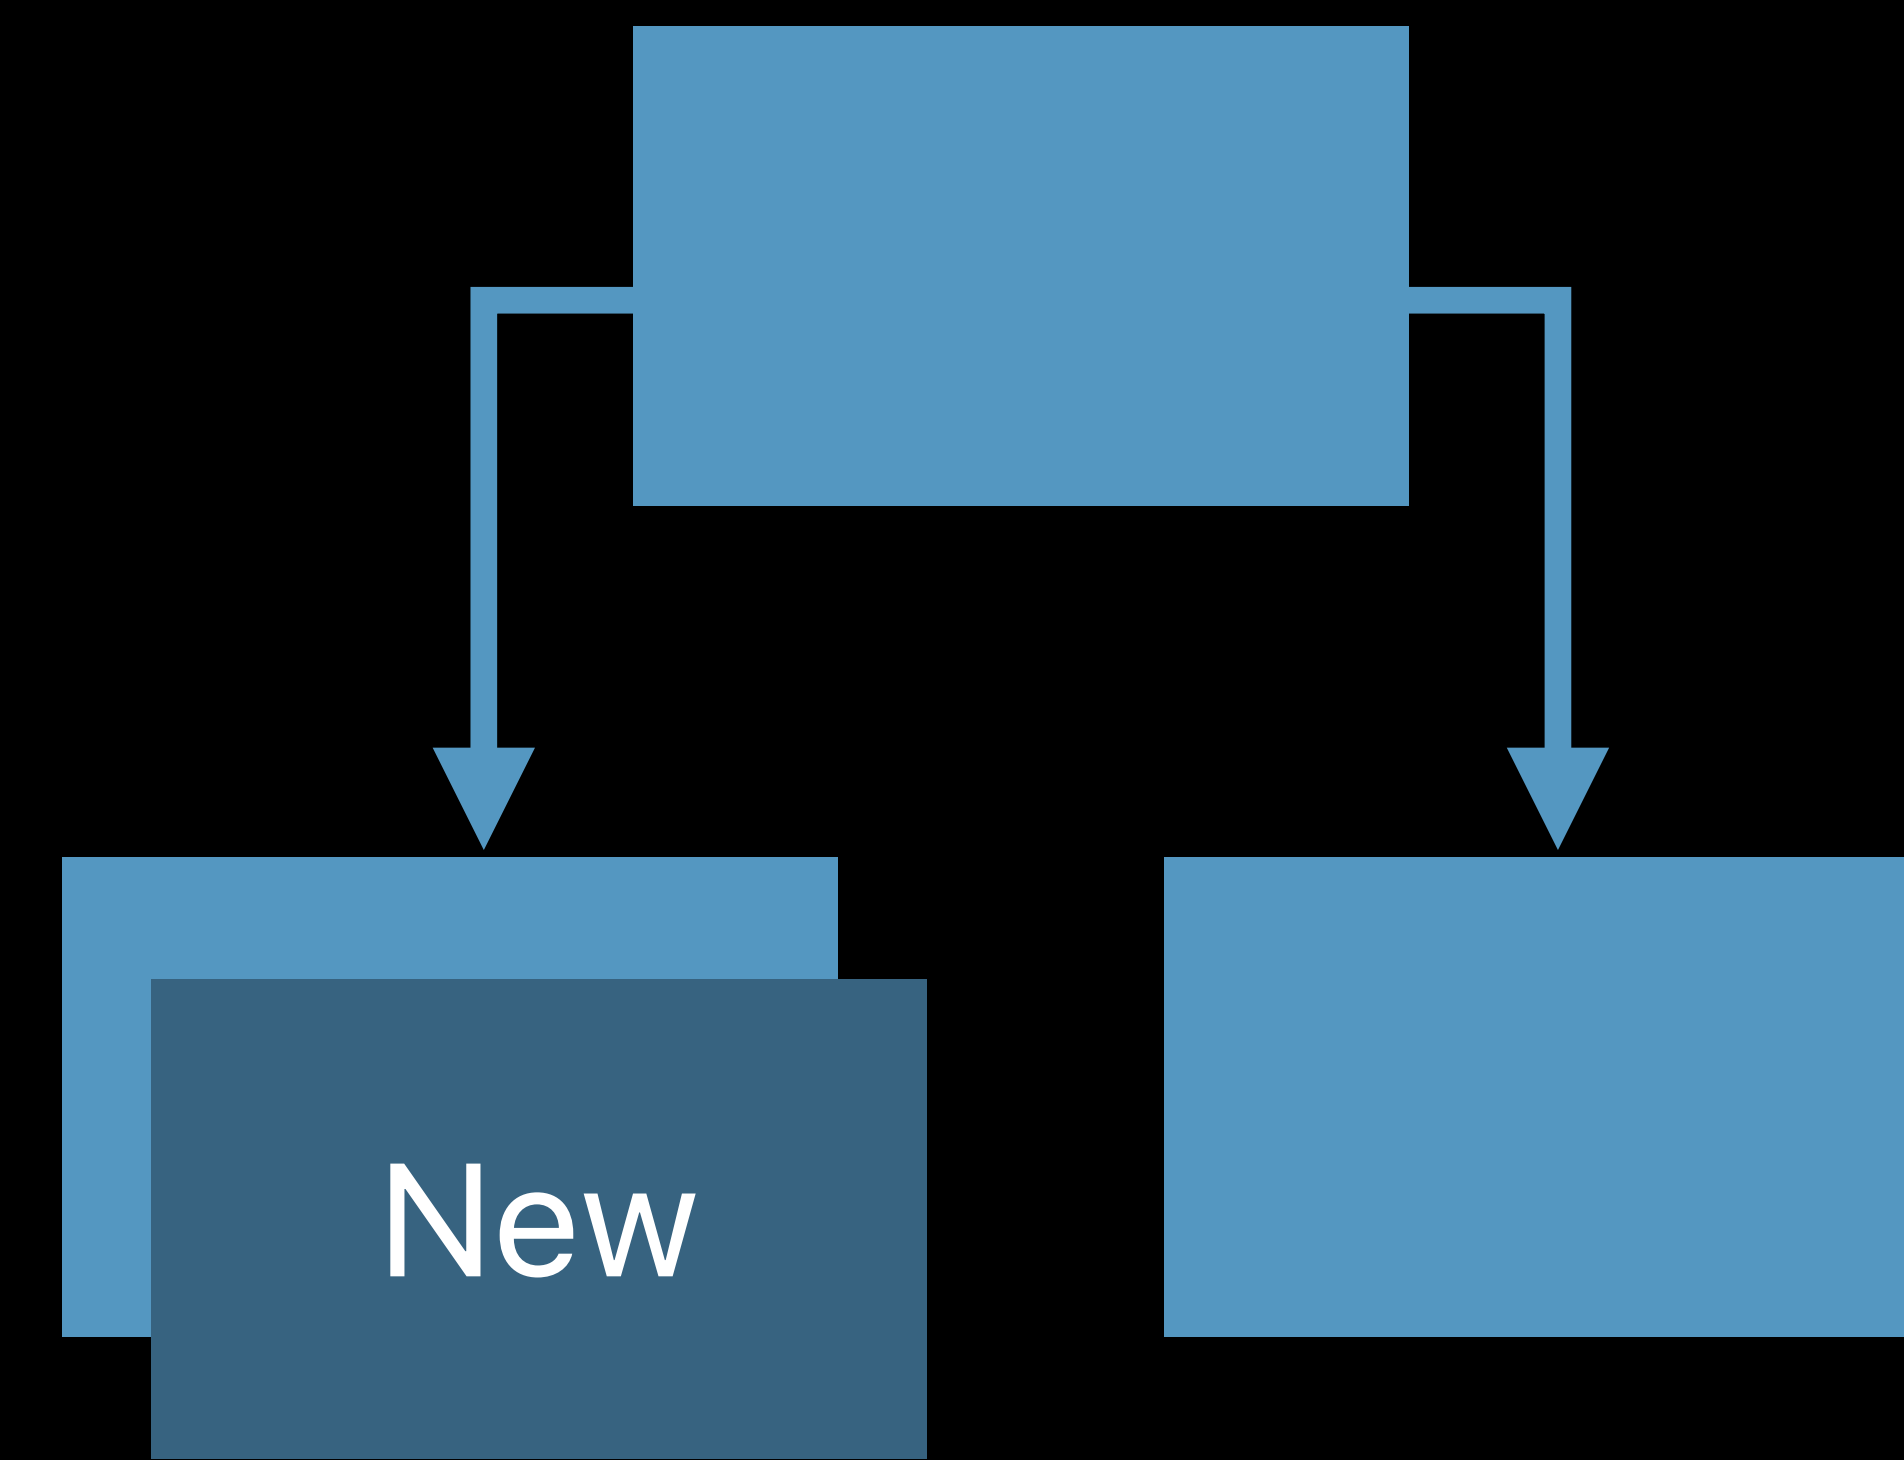




# File System Metadata

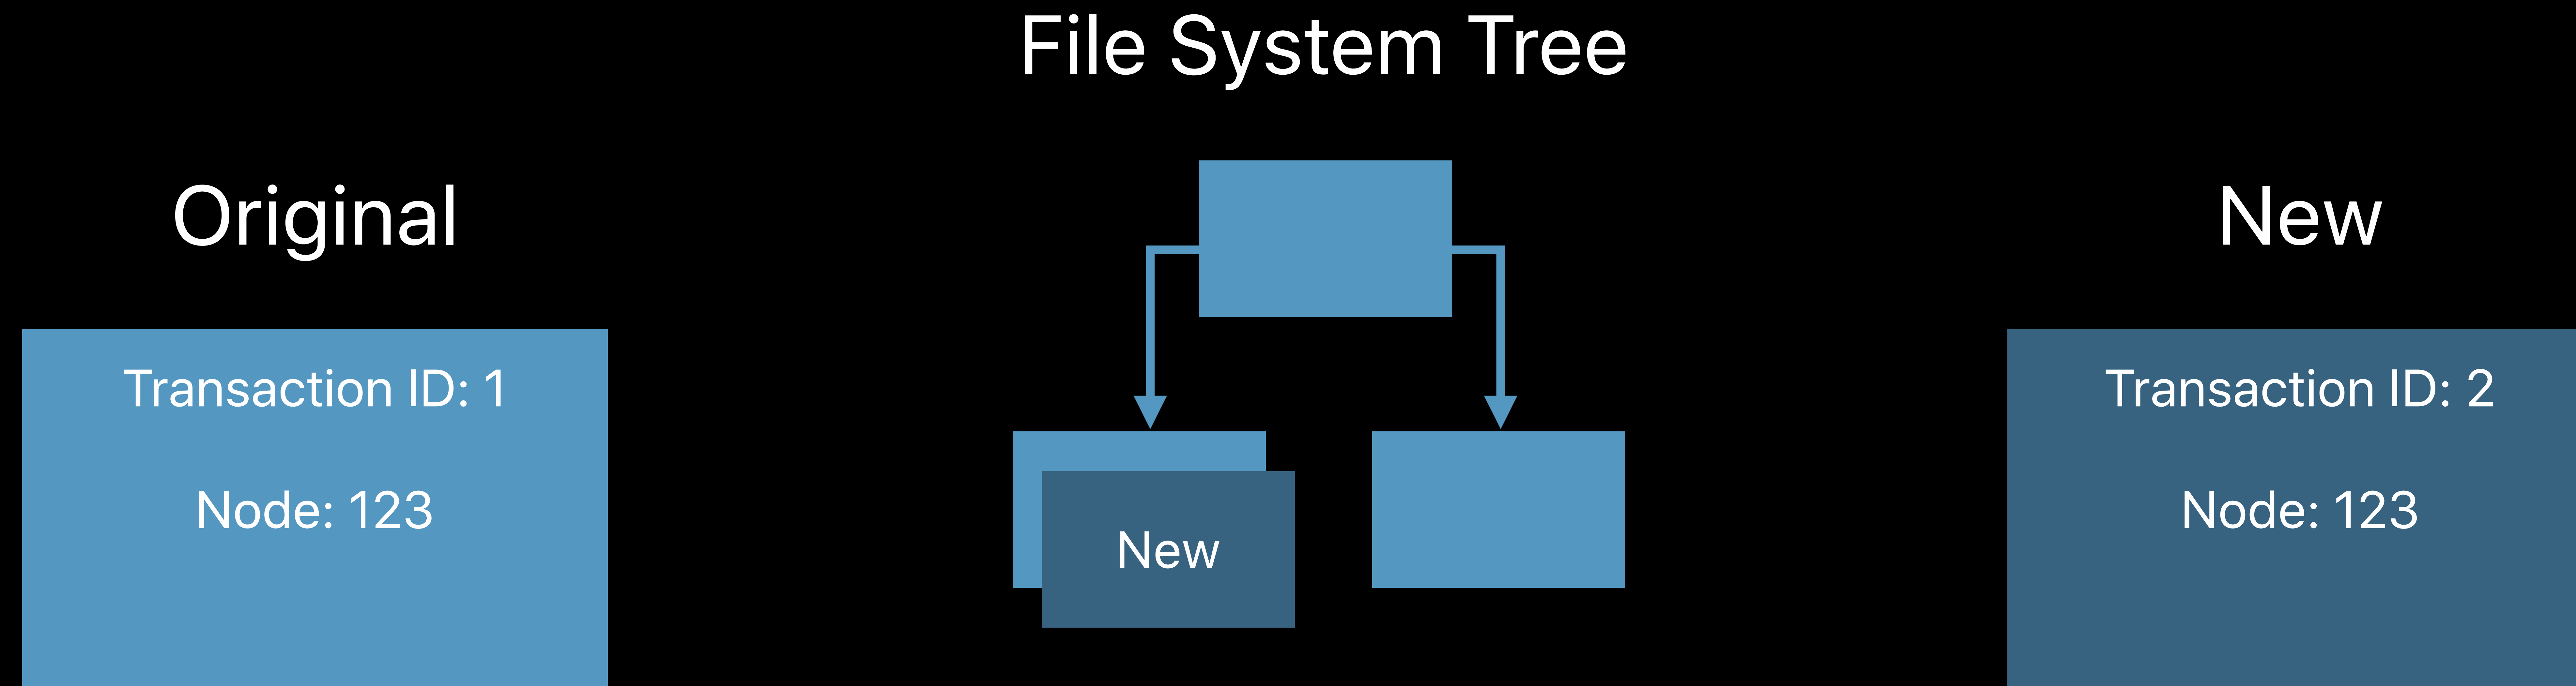
Writing a 240 byte NSDictionary to a file

## File System Tree



# File System Metadata

Writing a 240 byte NSDictionary to a file





# File System Metadata

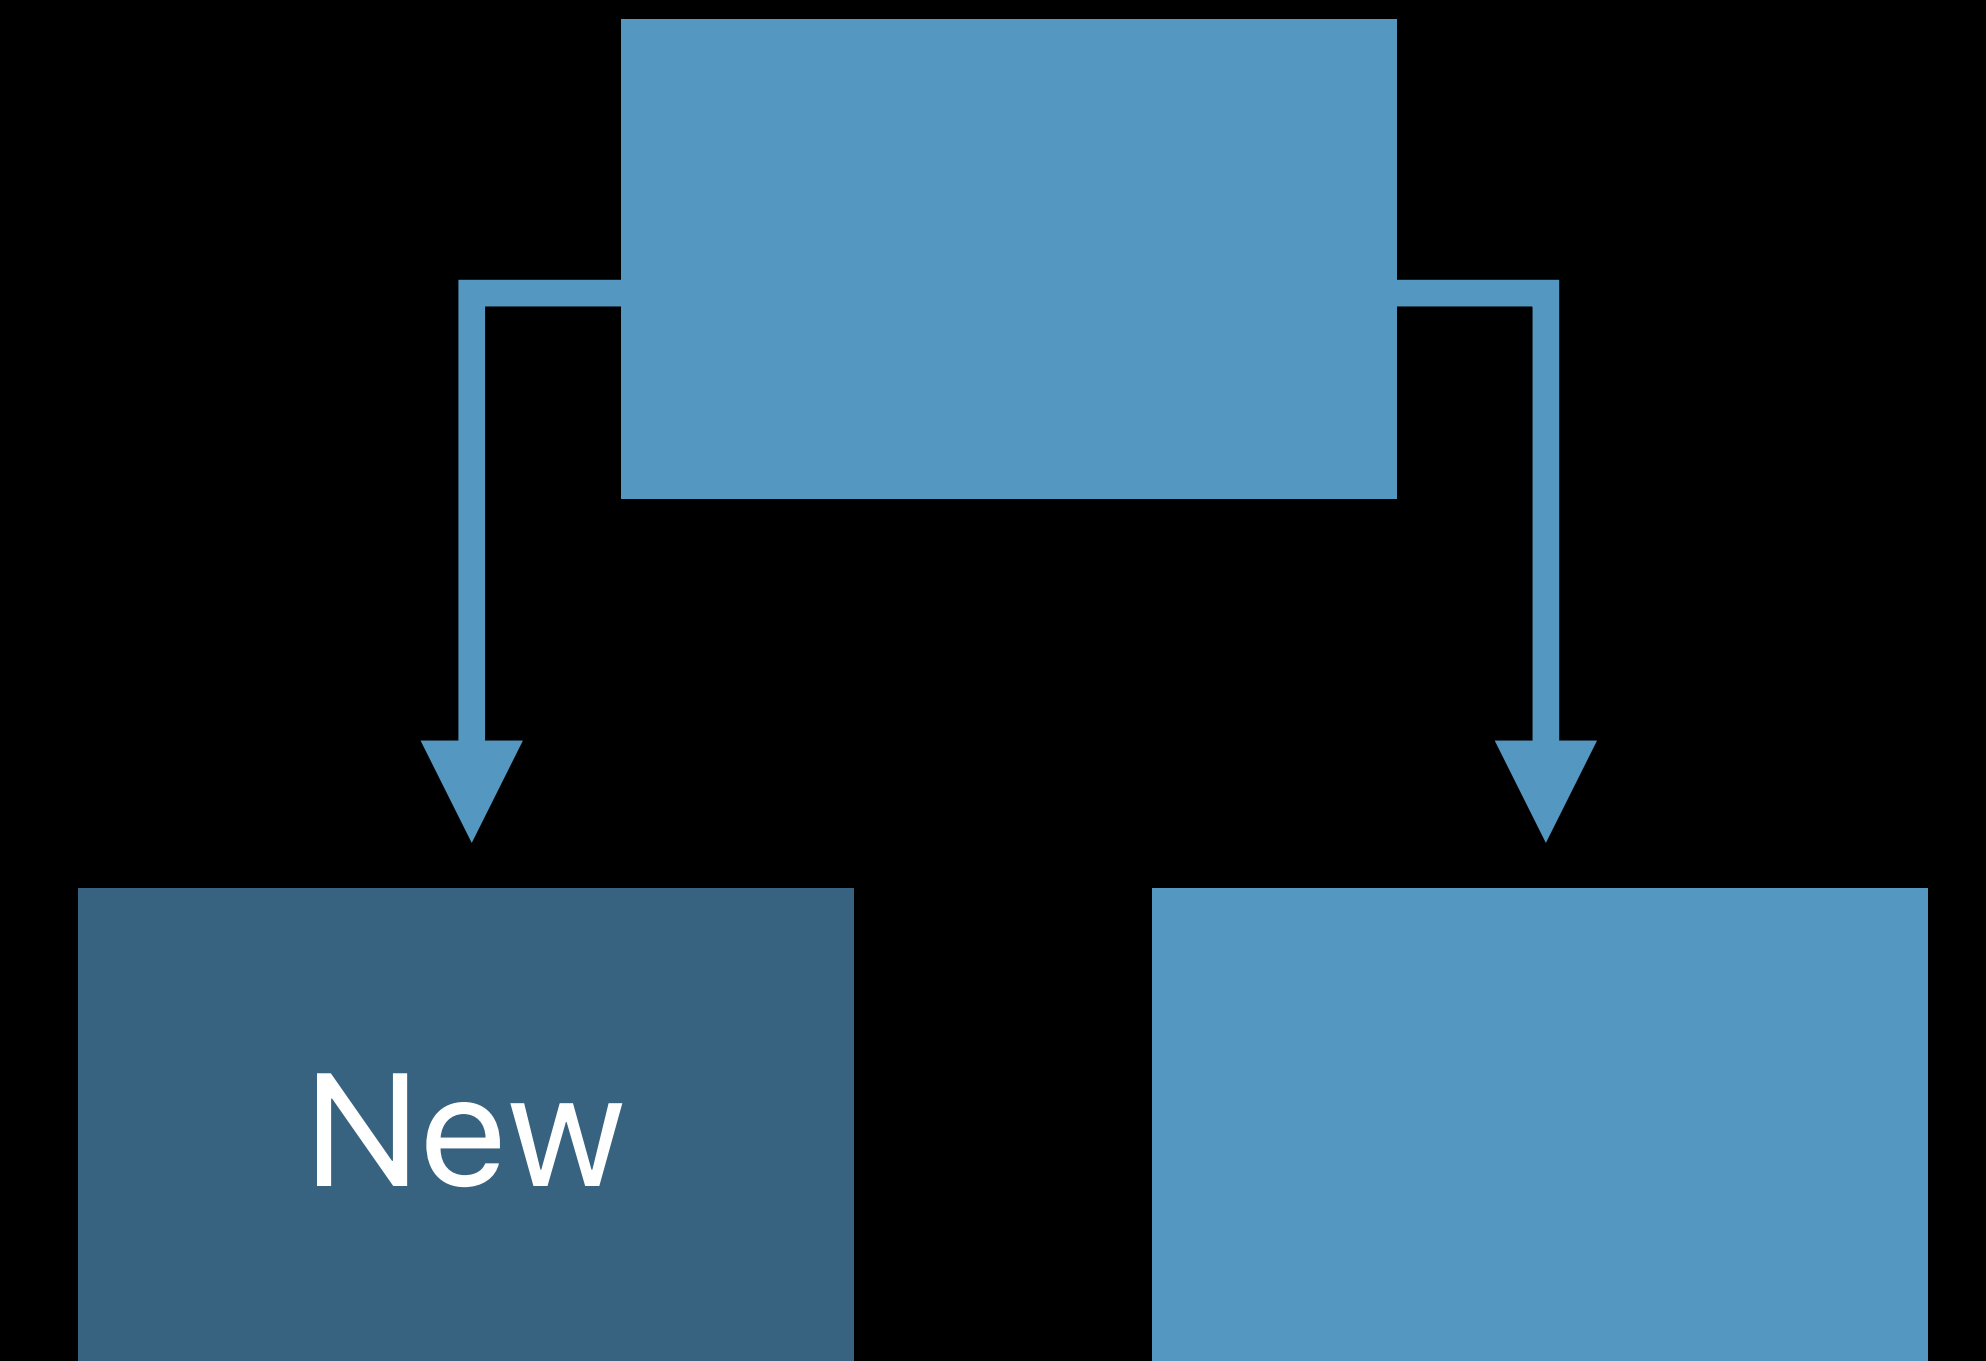
Writing a 240 byte NSDictionary to a file

Update file system node

Object Map



File System Tree



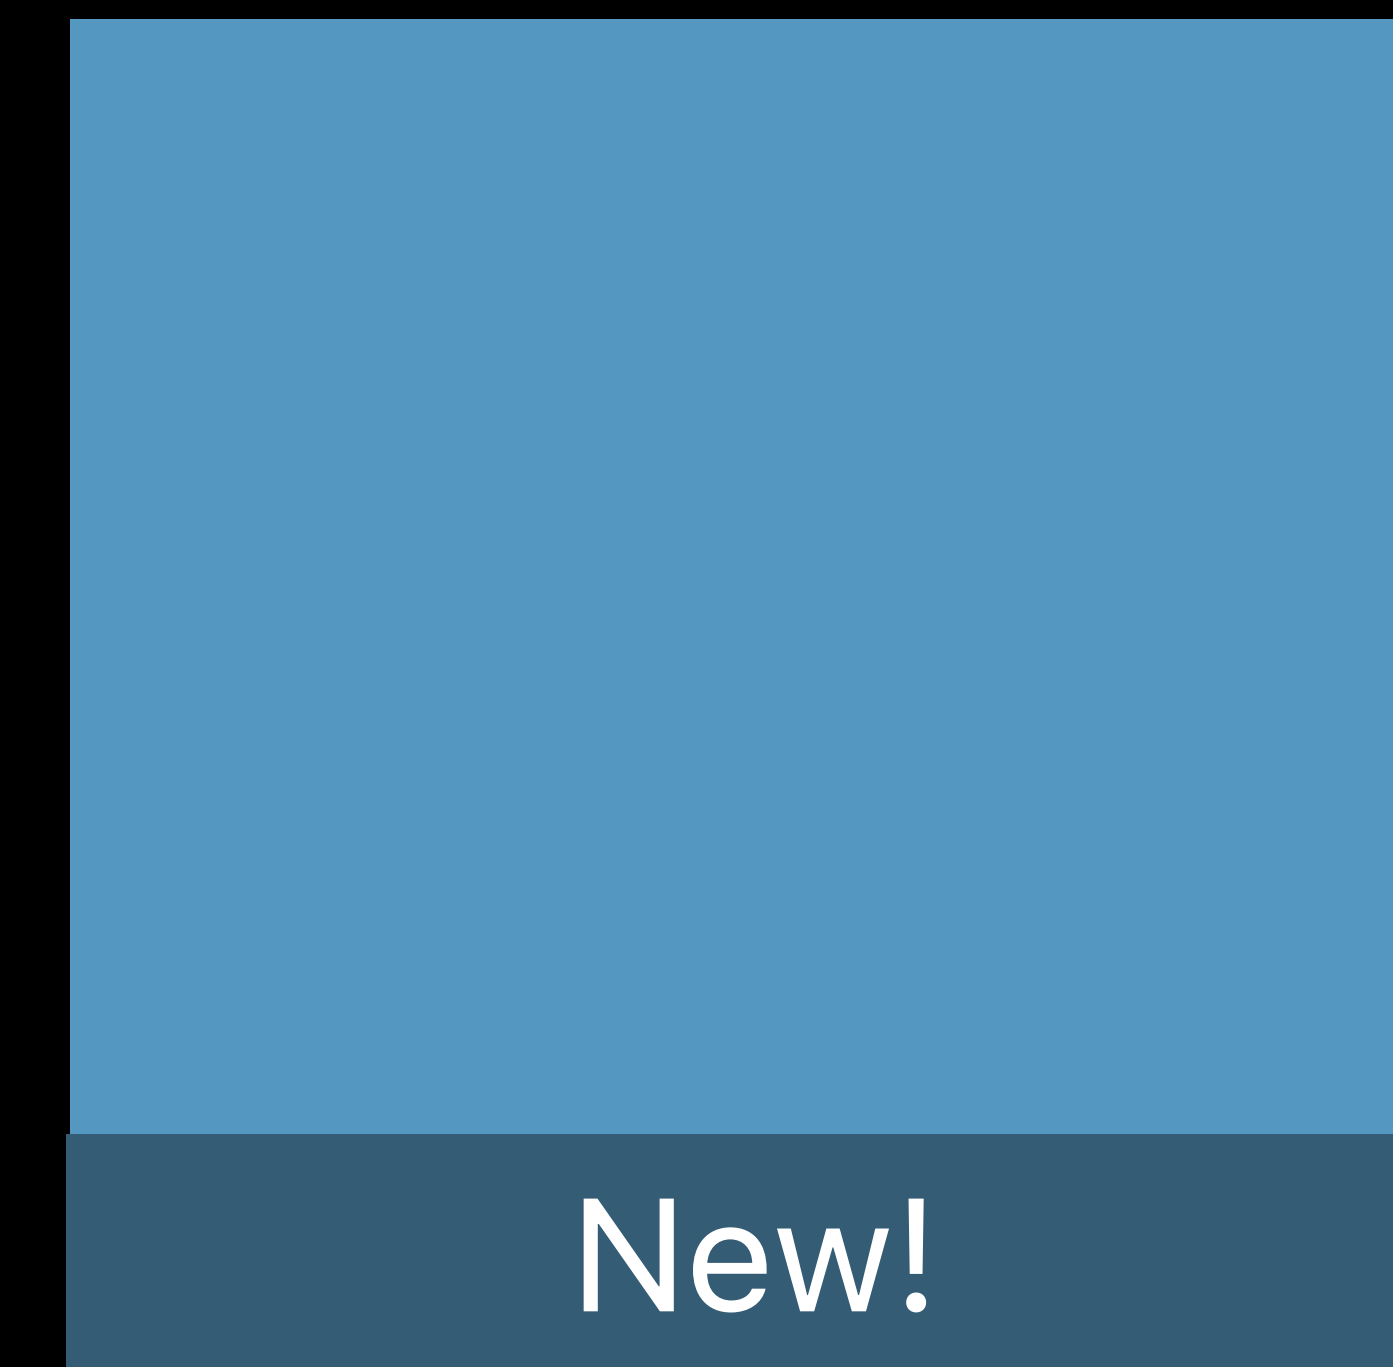
# File System Metadata

Writing a 240 byte NSDictionary to a file

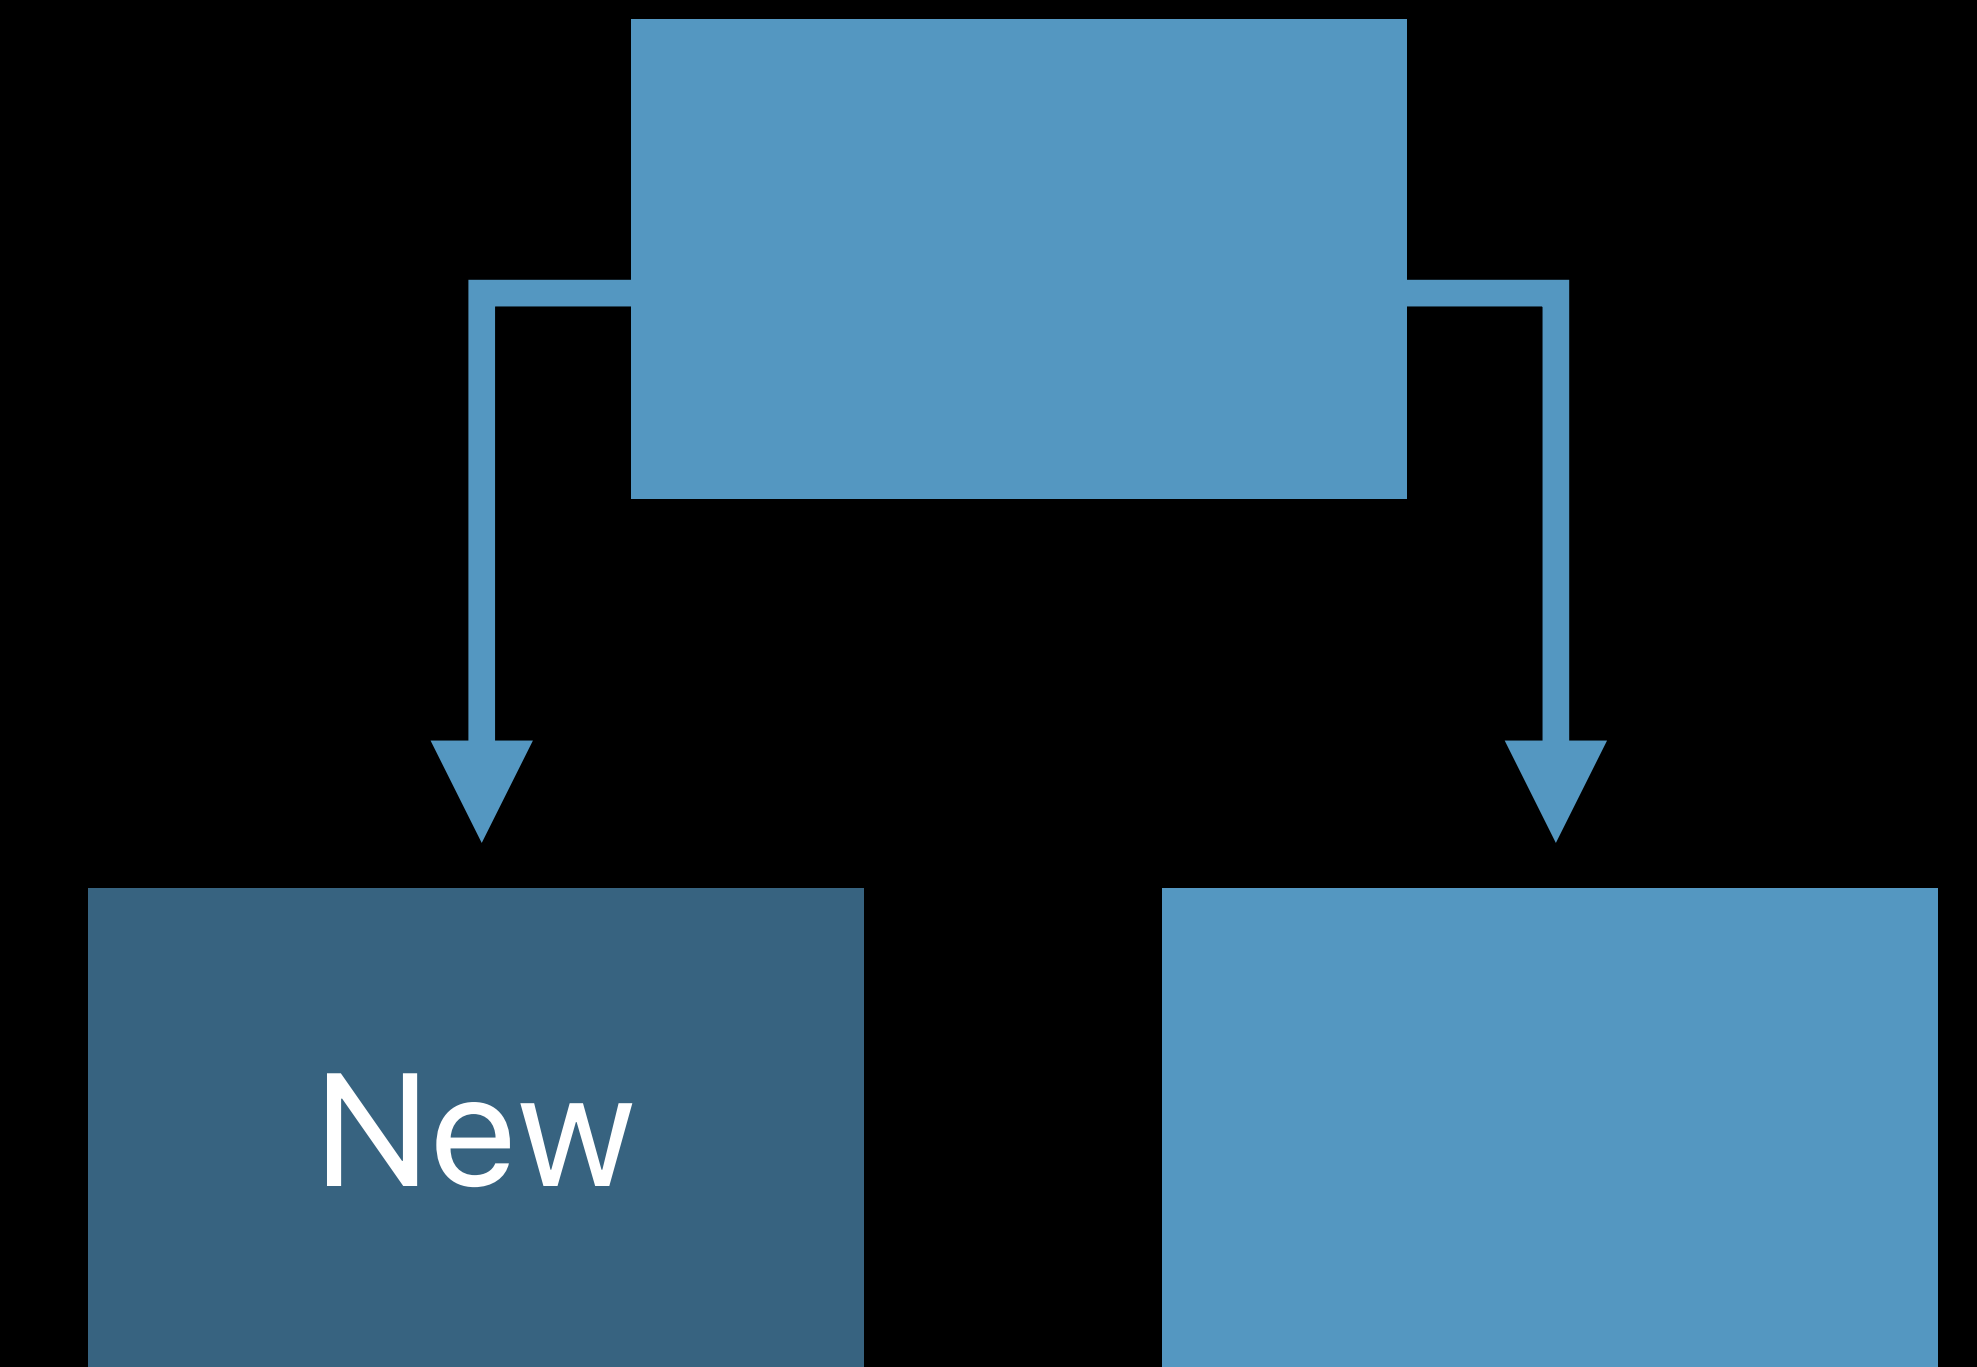
Update file system node

Update object map

Object Map



File System Tree





# File System Metadata

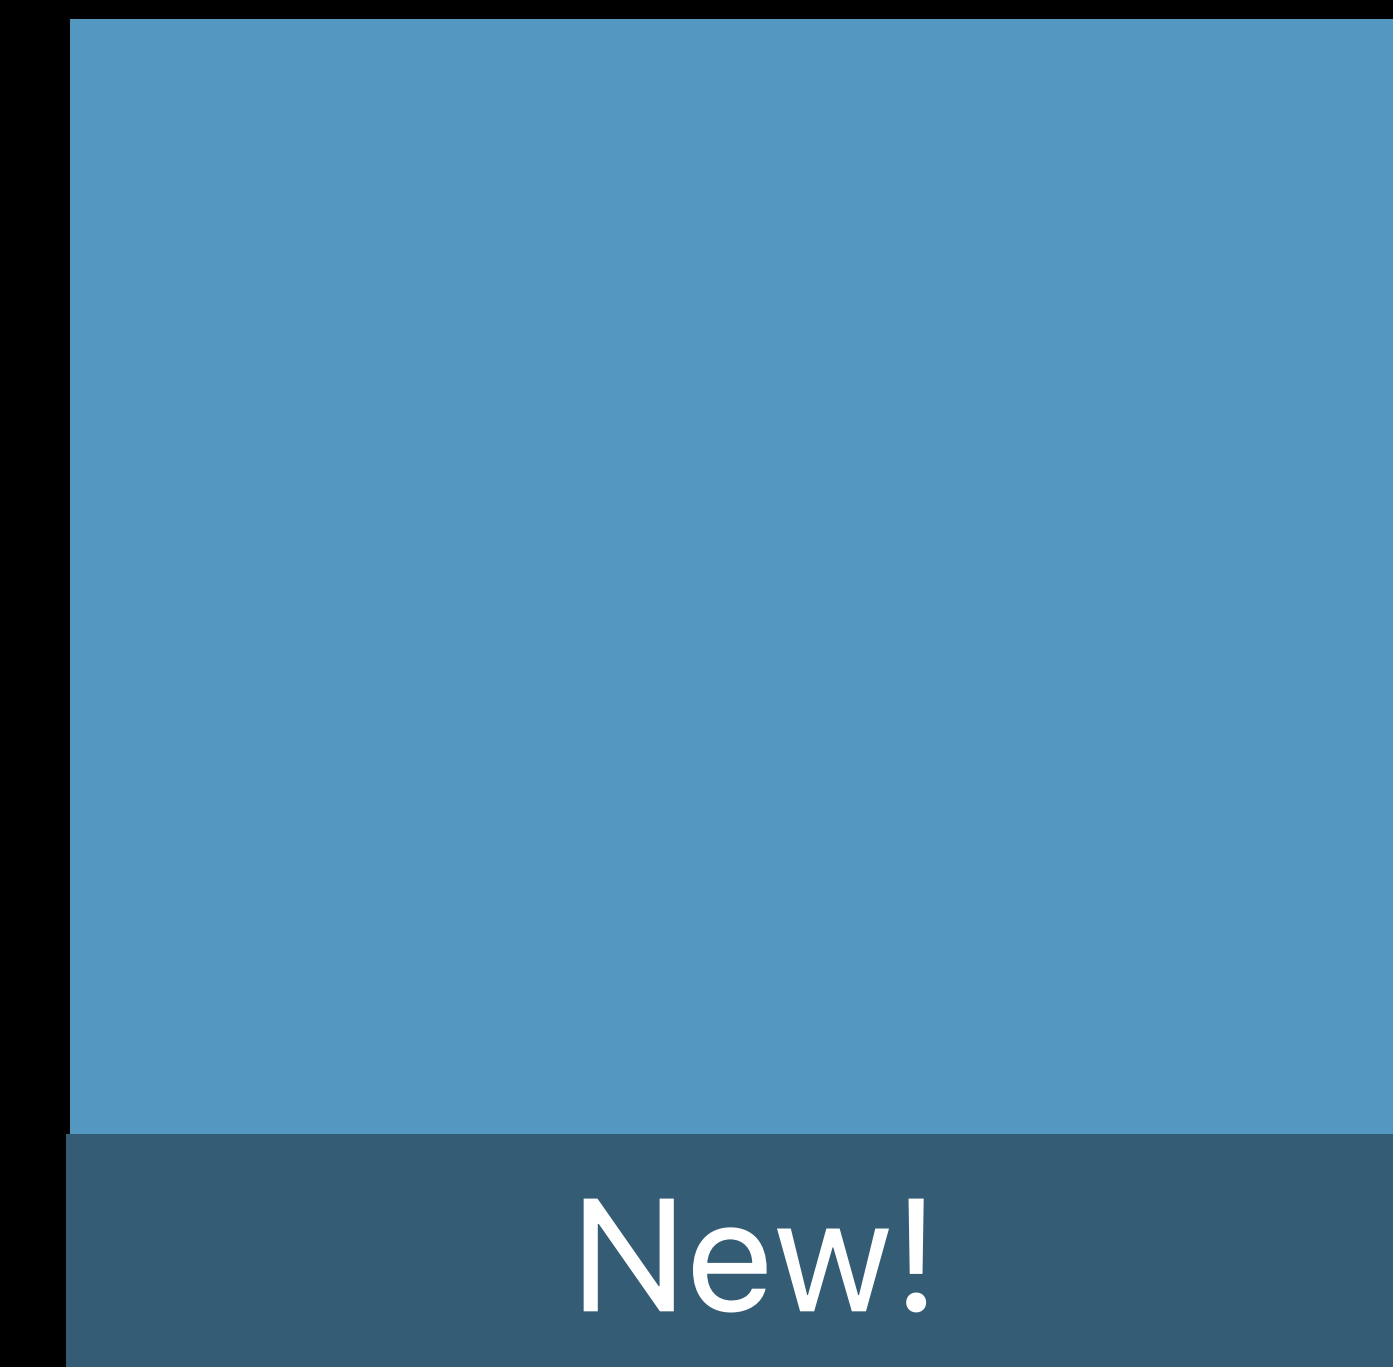
Writing a 240 byte NSDictionary to a file

Update file system node

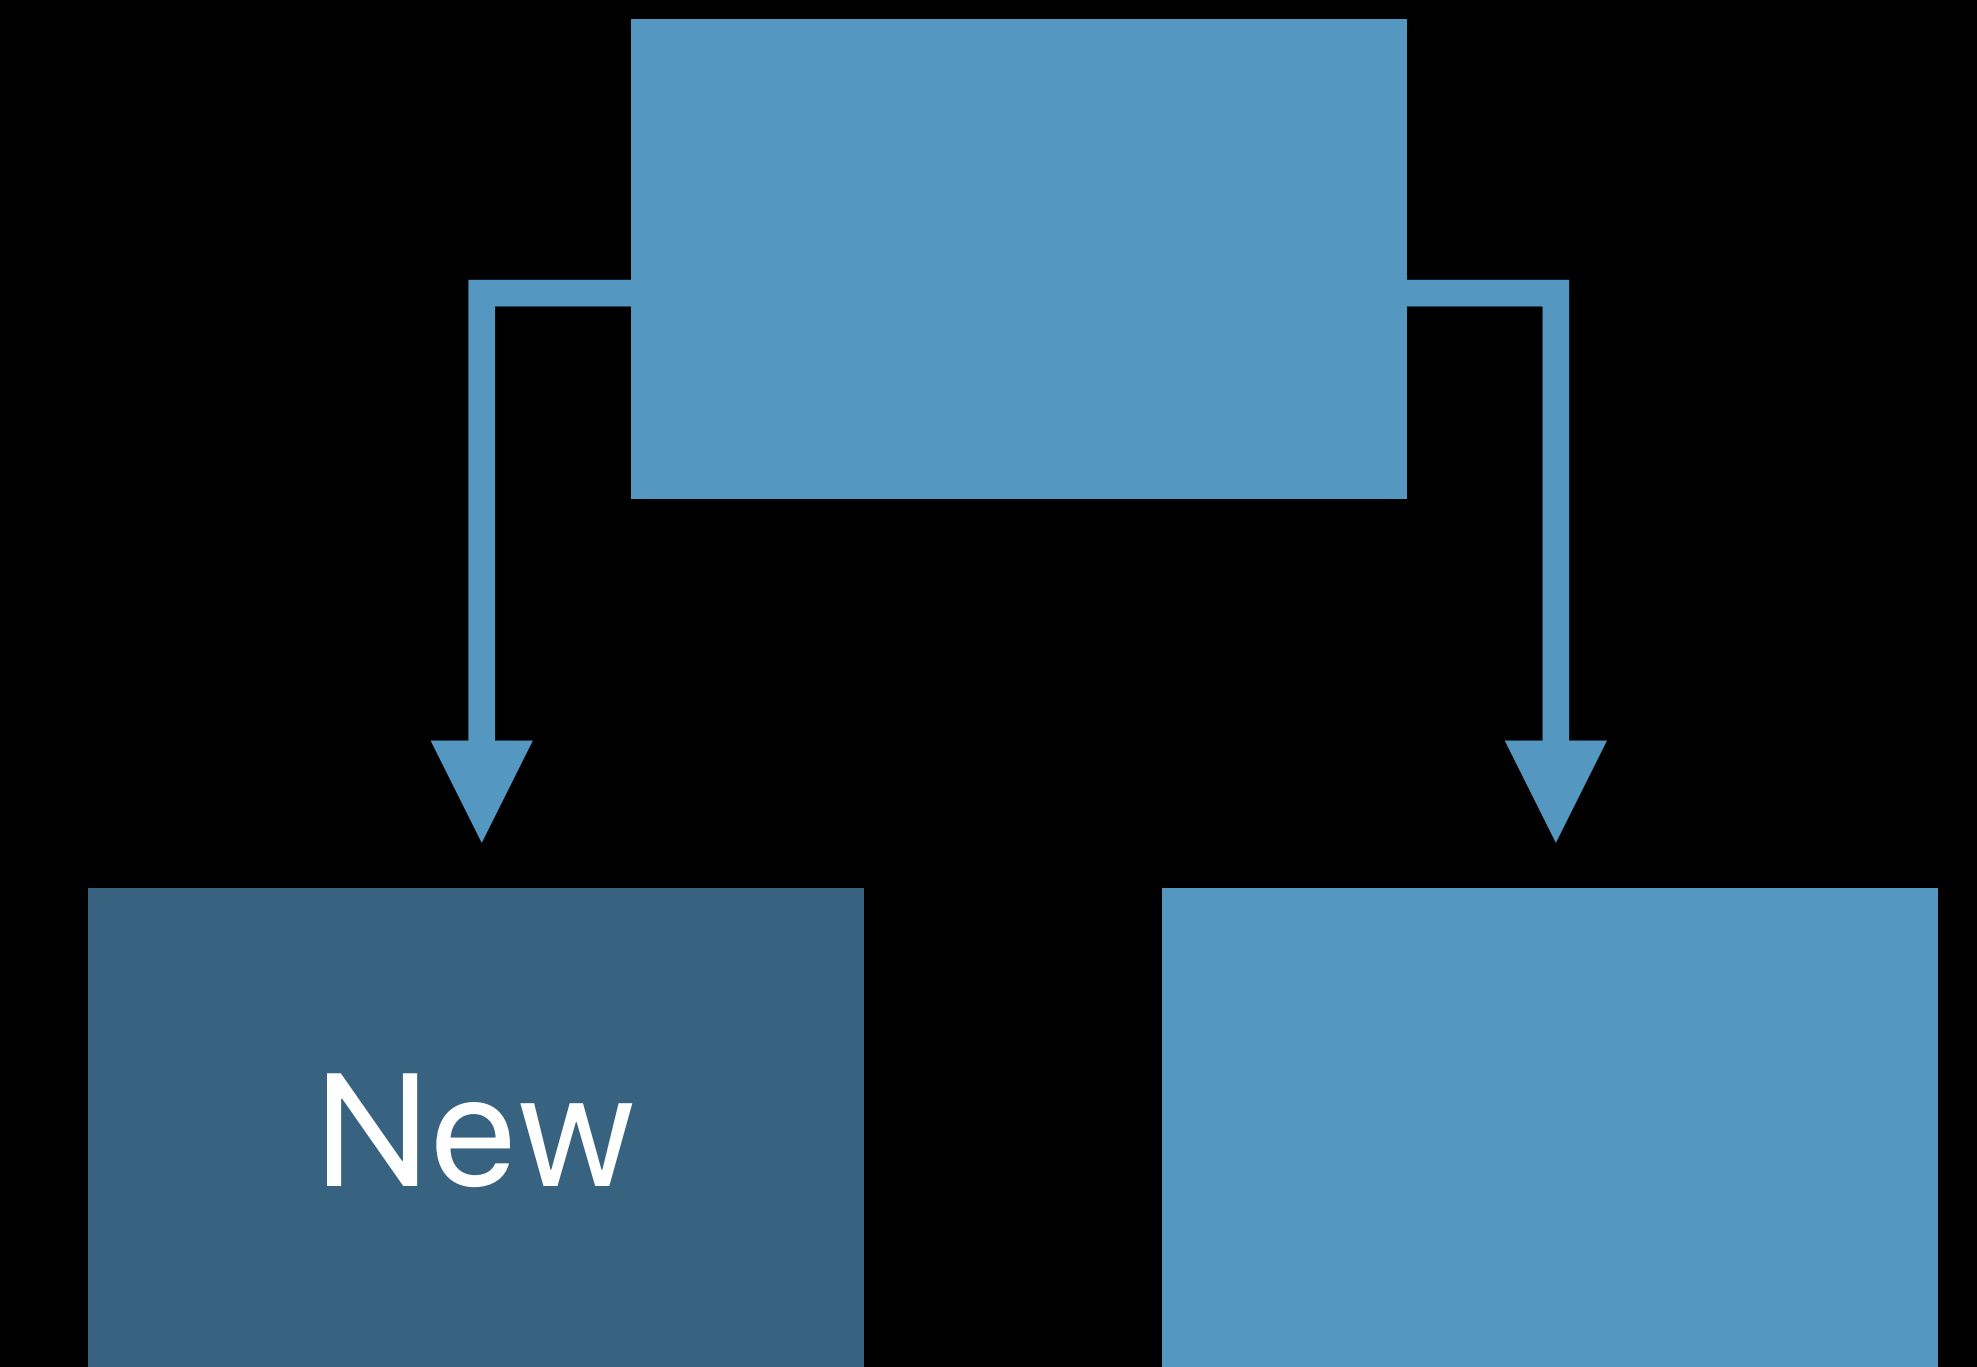
Update object map

Total Metadata: 8K

Object Map



File System Tree



# File System Metadata

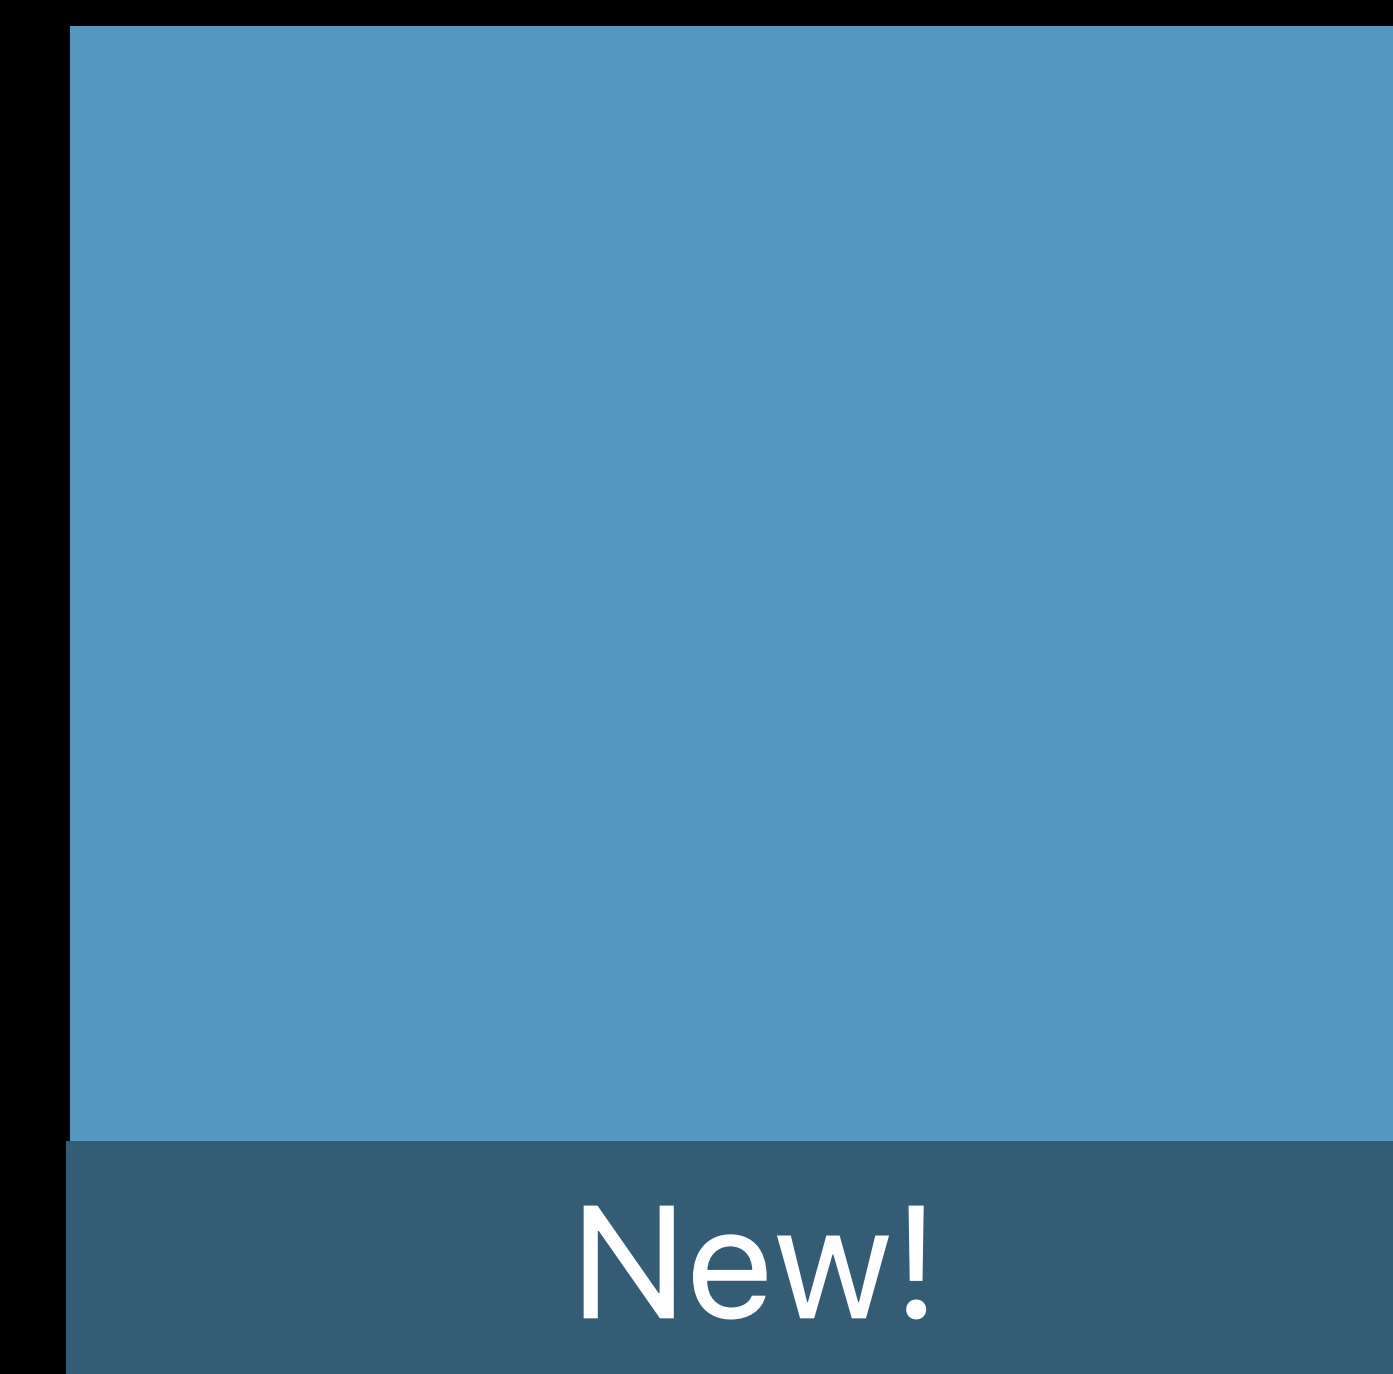
Writing a 240 byte NSDictionary to a file

Update file system node (4K)

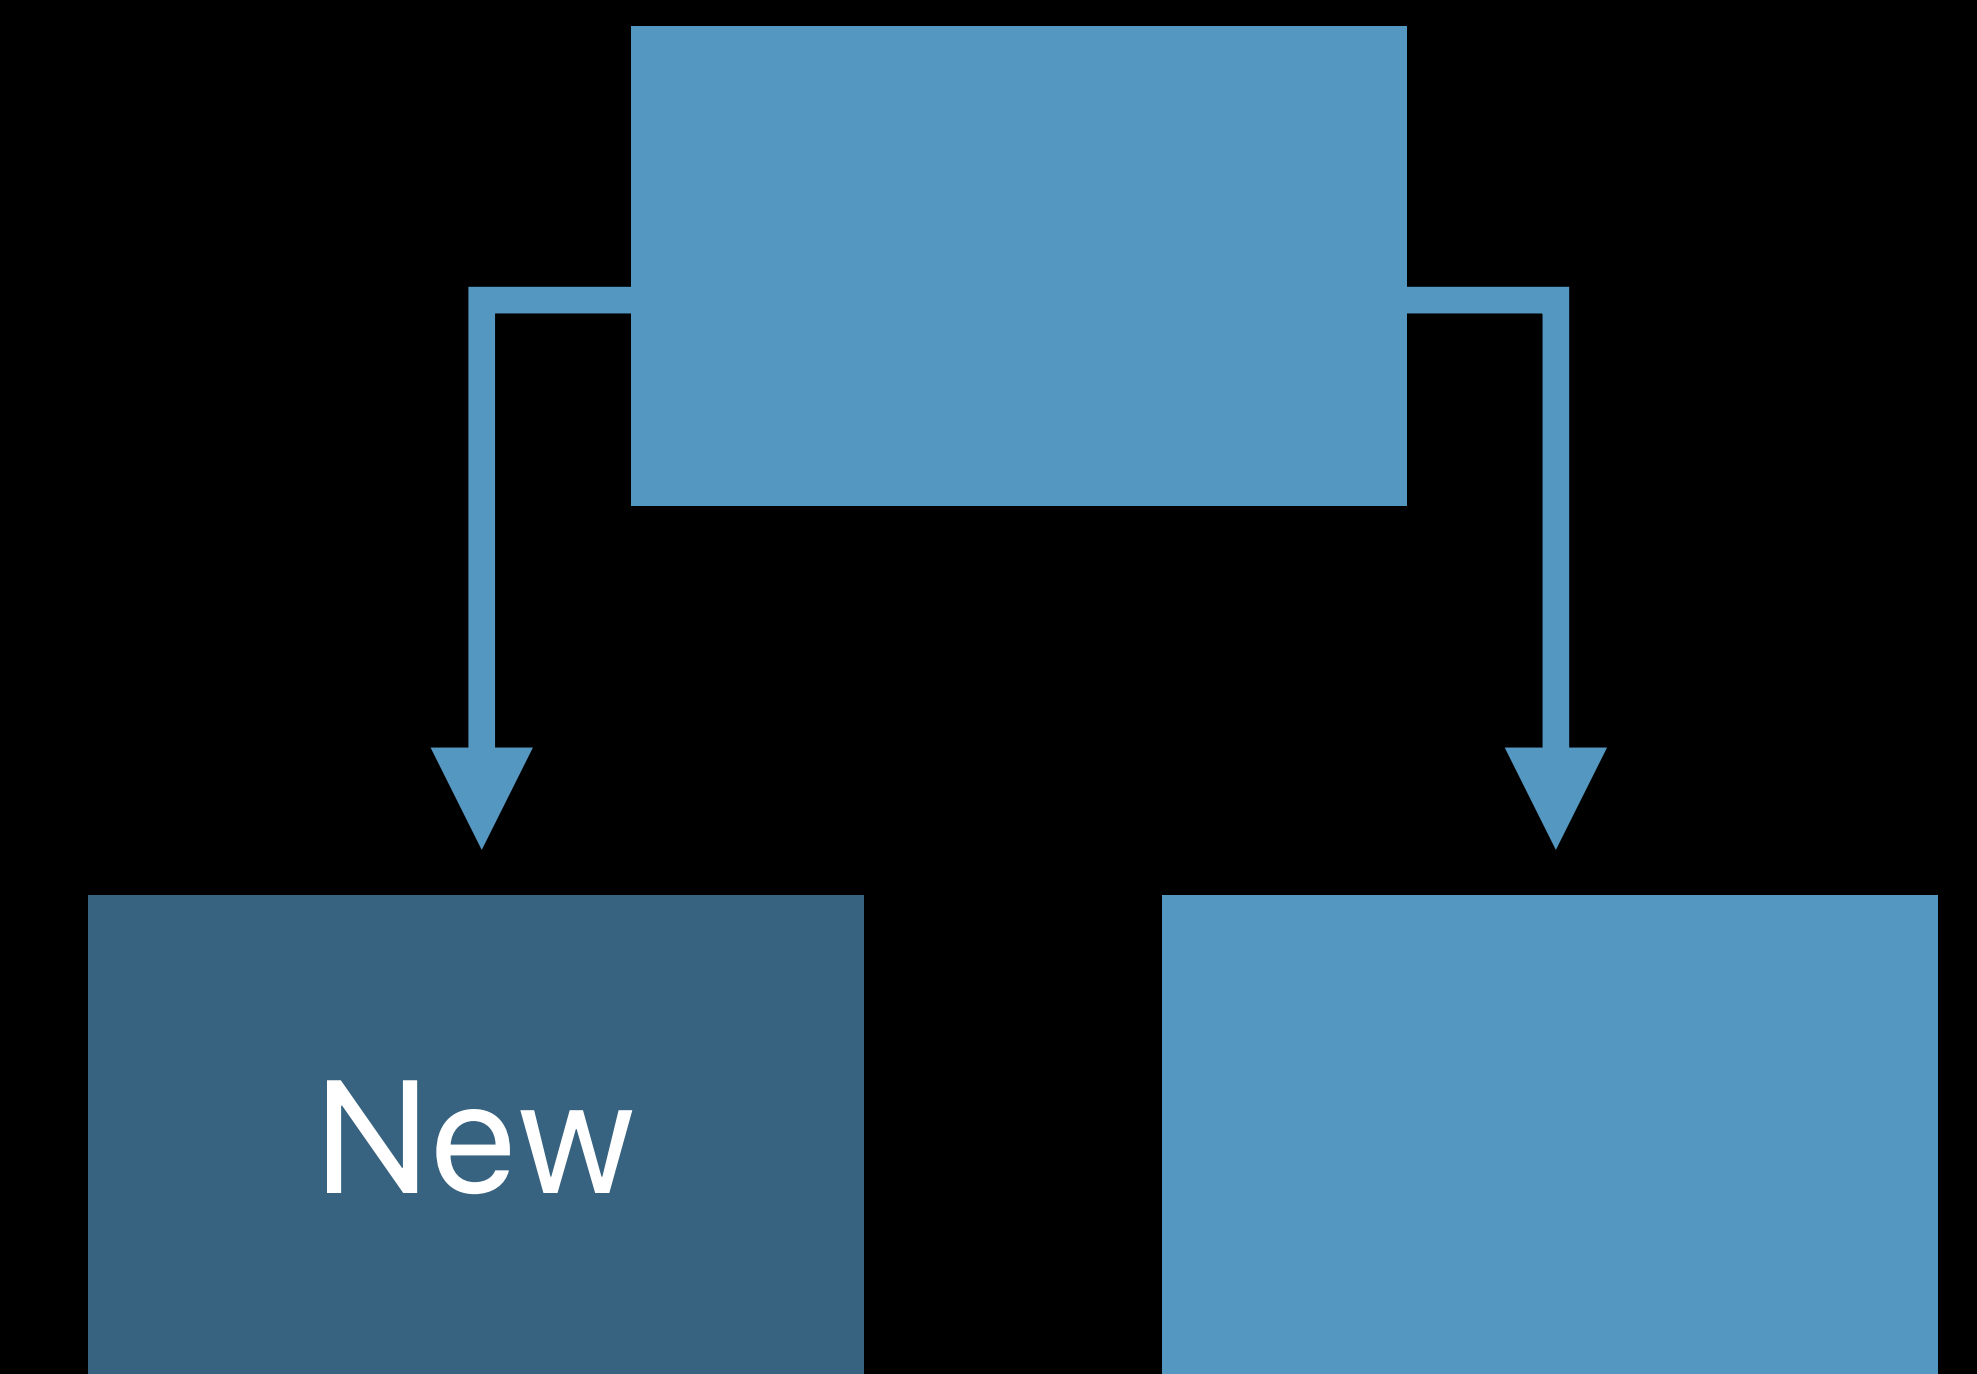
Update object map

Total Metadata: 8K

Object Map



File System Tree





# File System Metadata

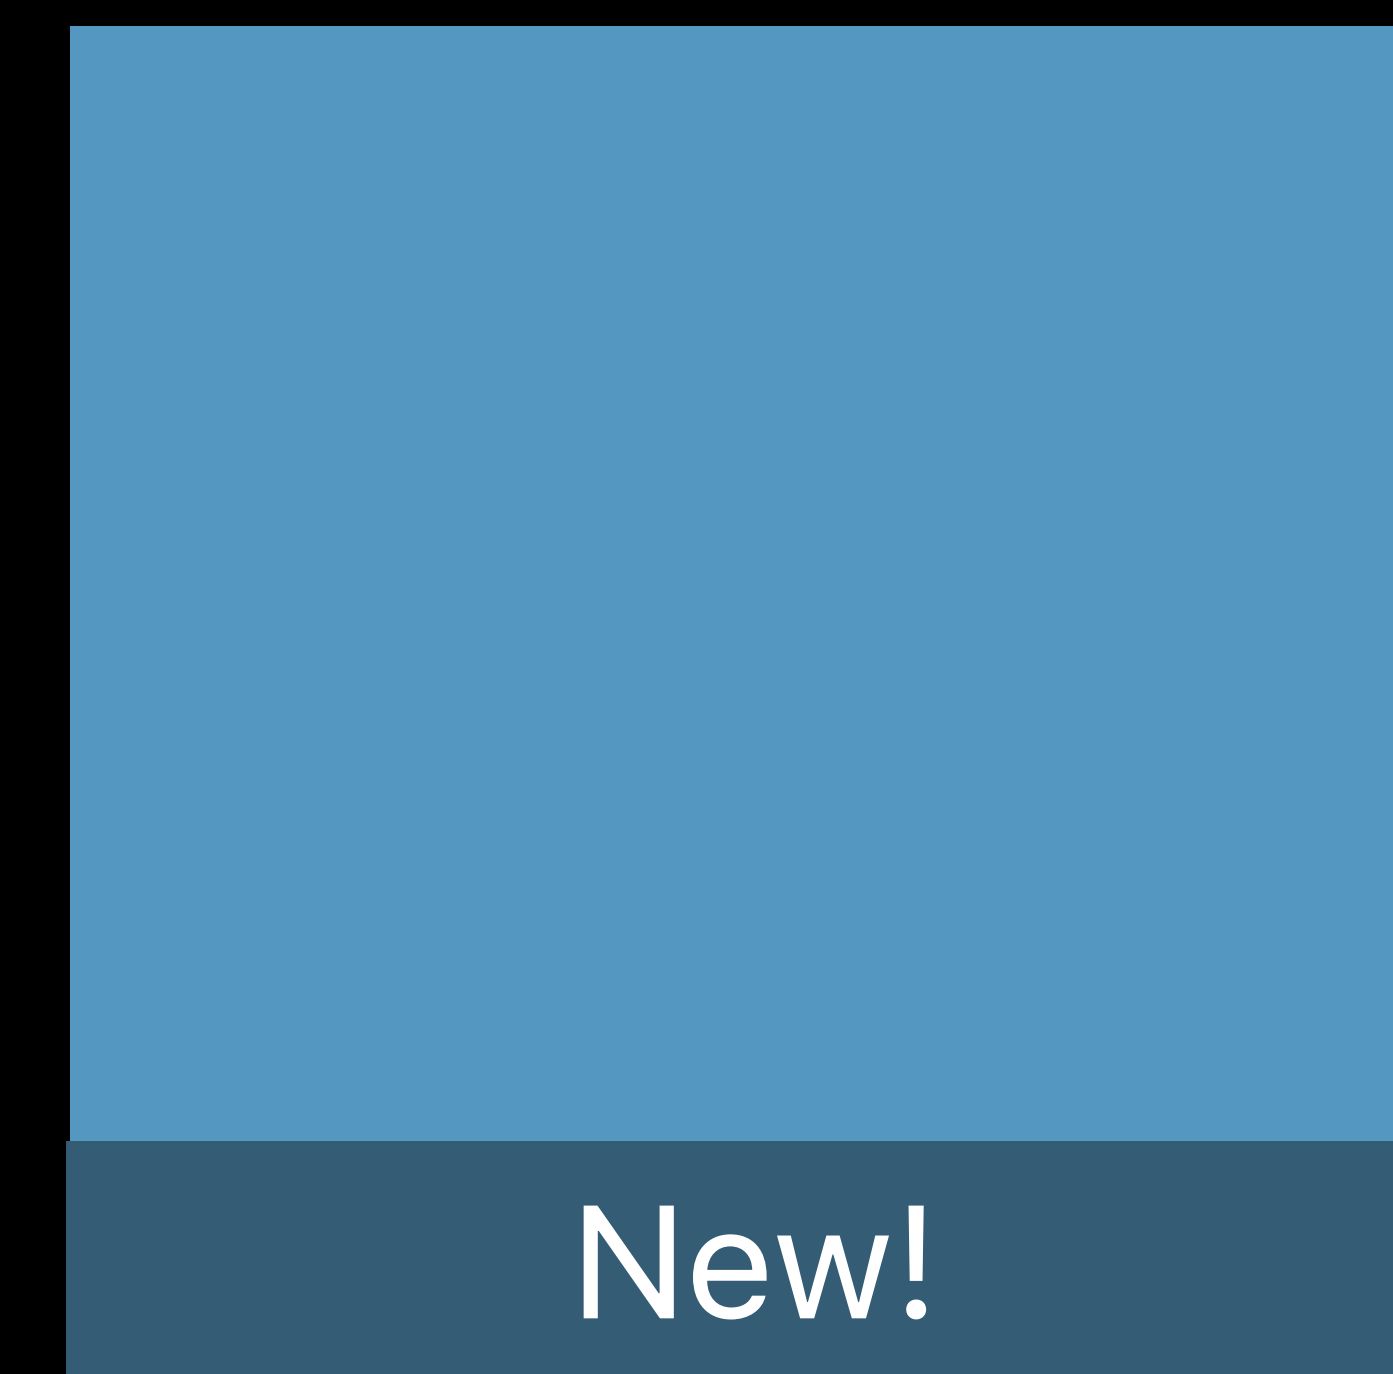
Writing a 240 byte NSDictionary to a file

Update file system node (4K)

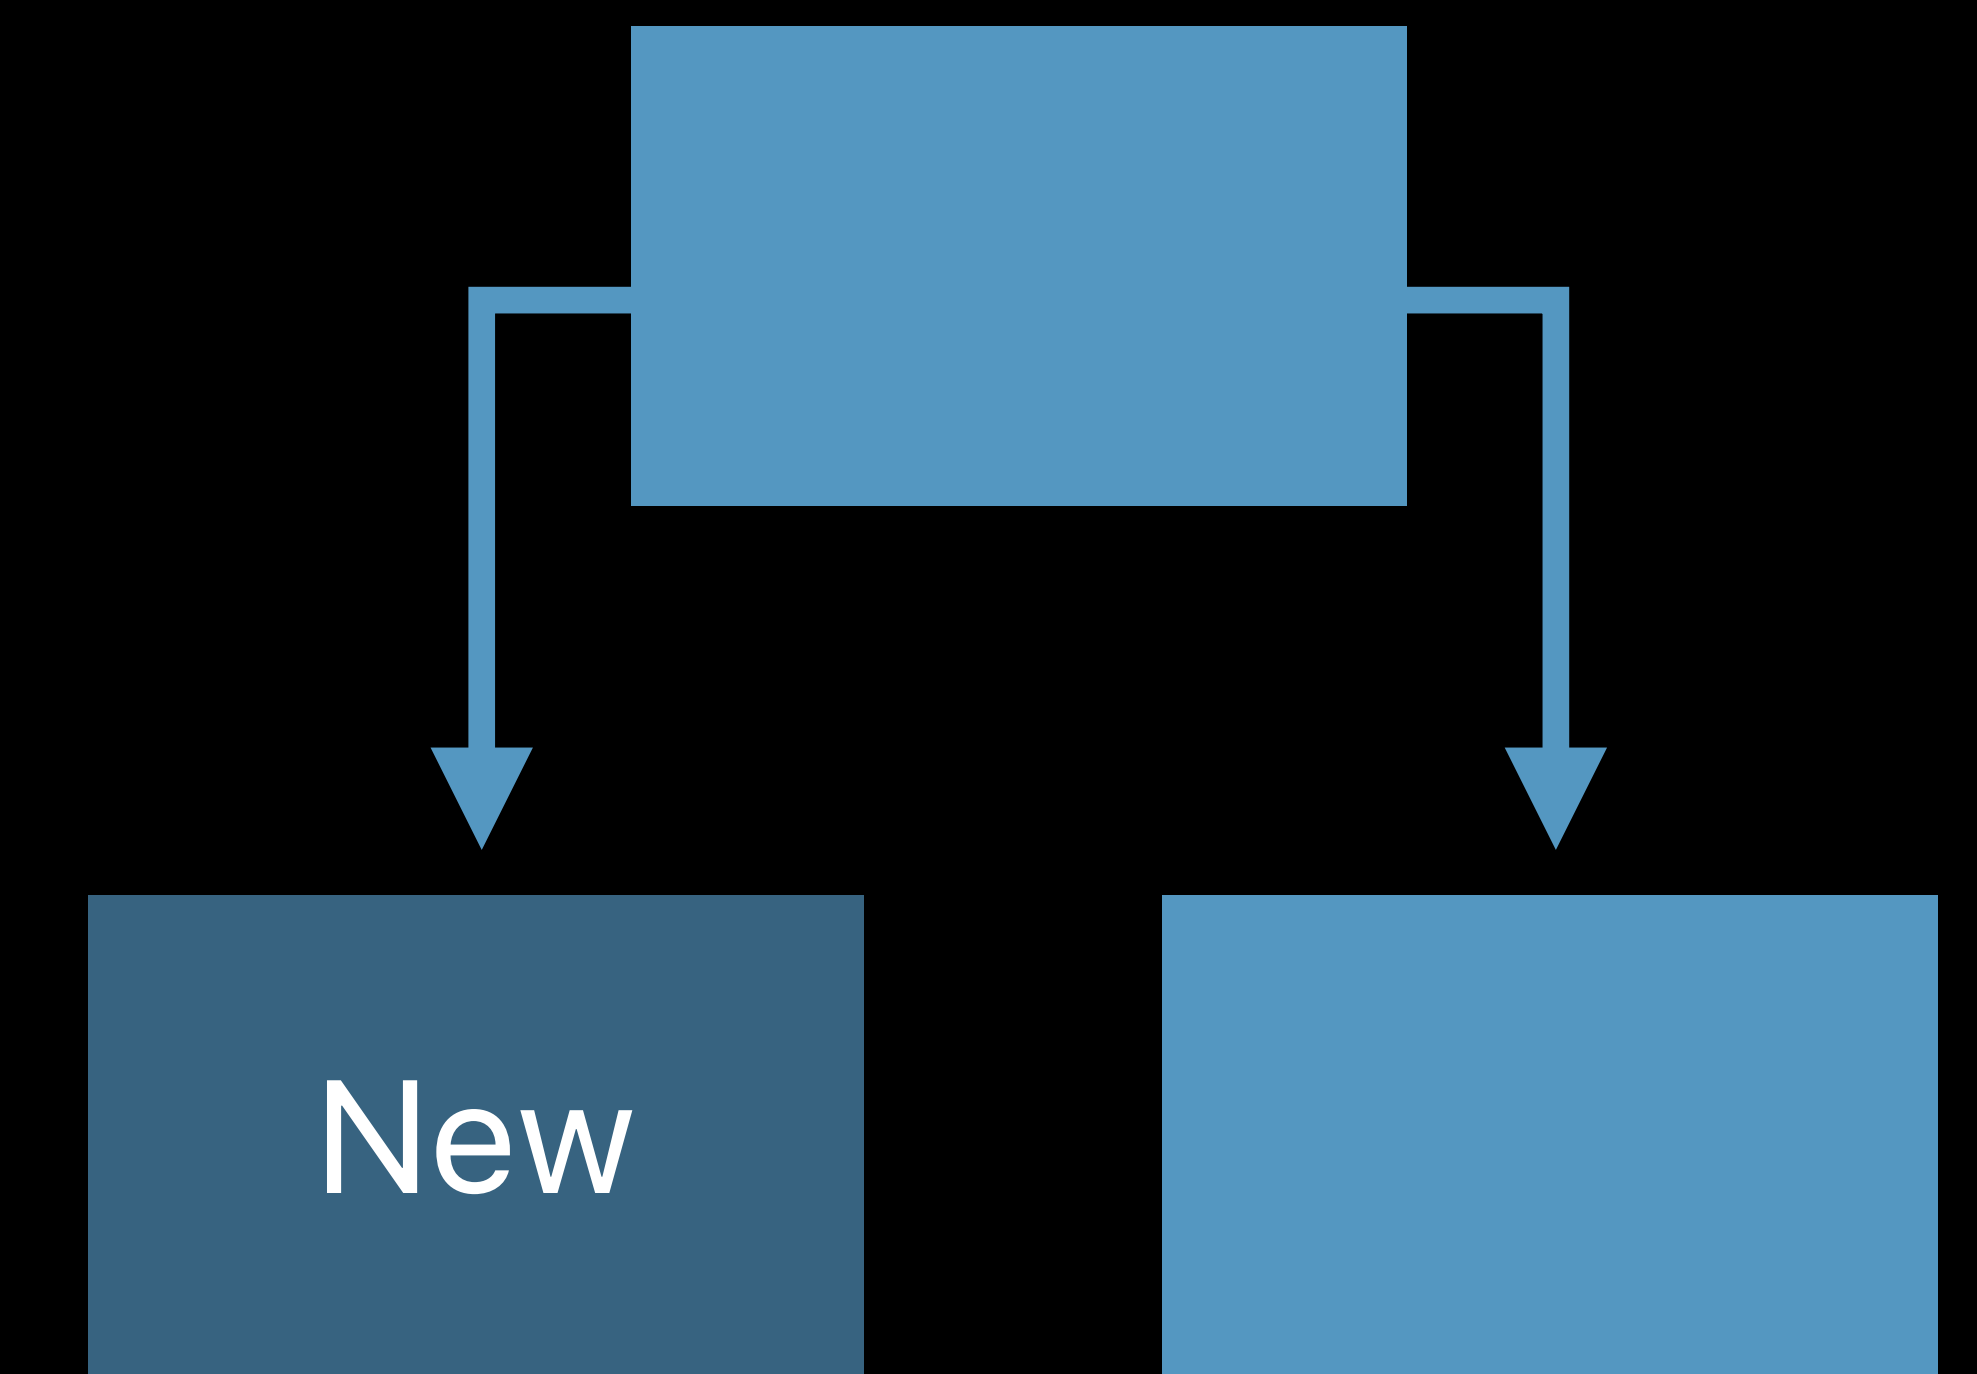
Update object map (4K)

Total Metadata: 8K

Object Map



File System Tree



# File System Metadata

Writing a 240 byte NSDictionary to a file

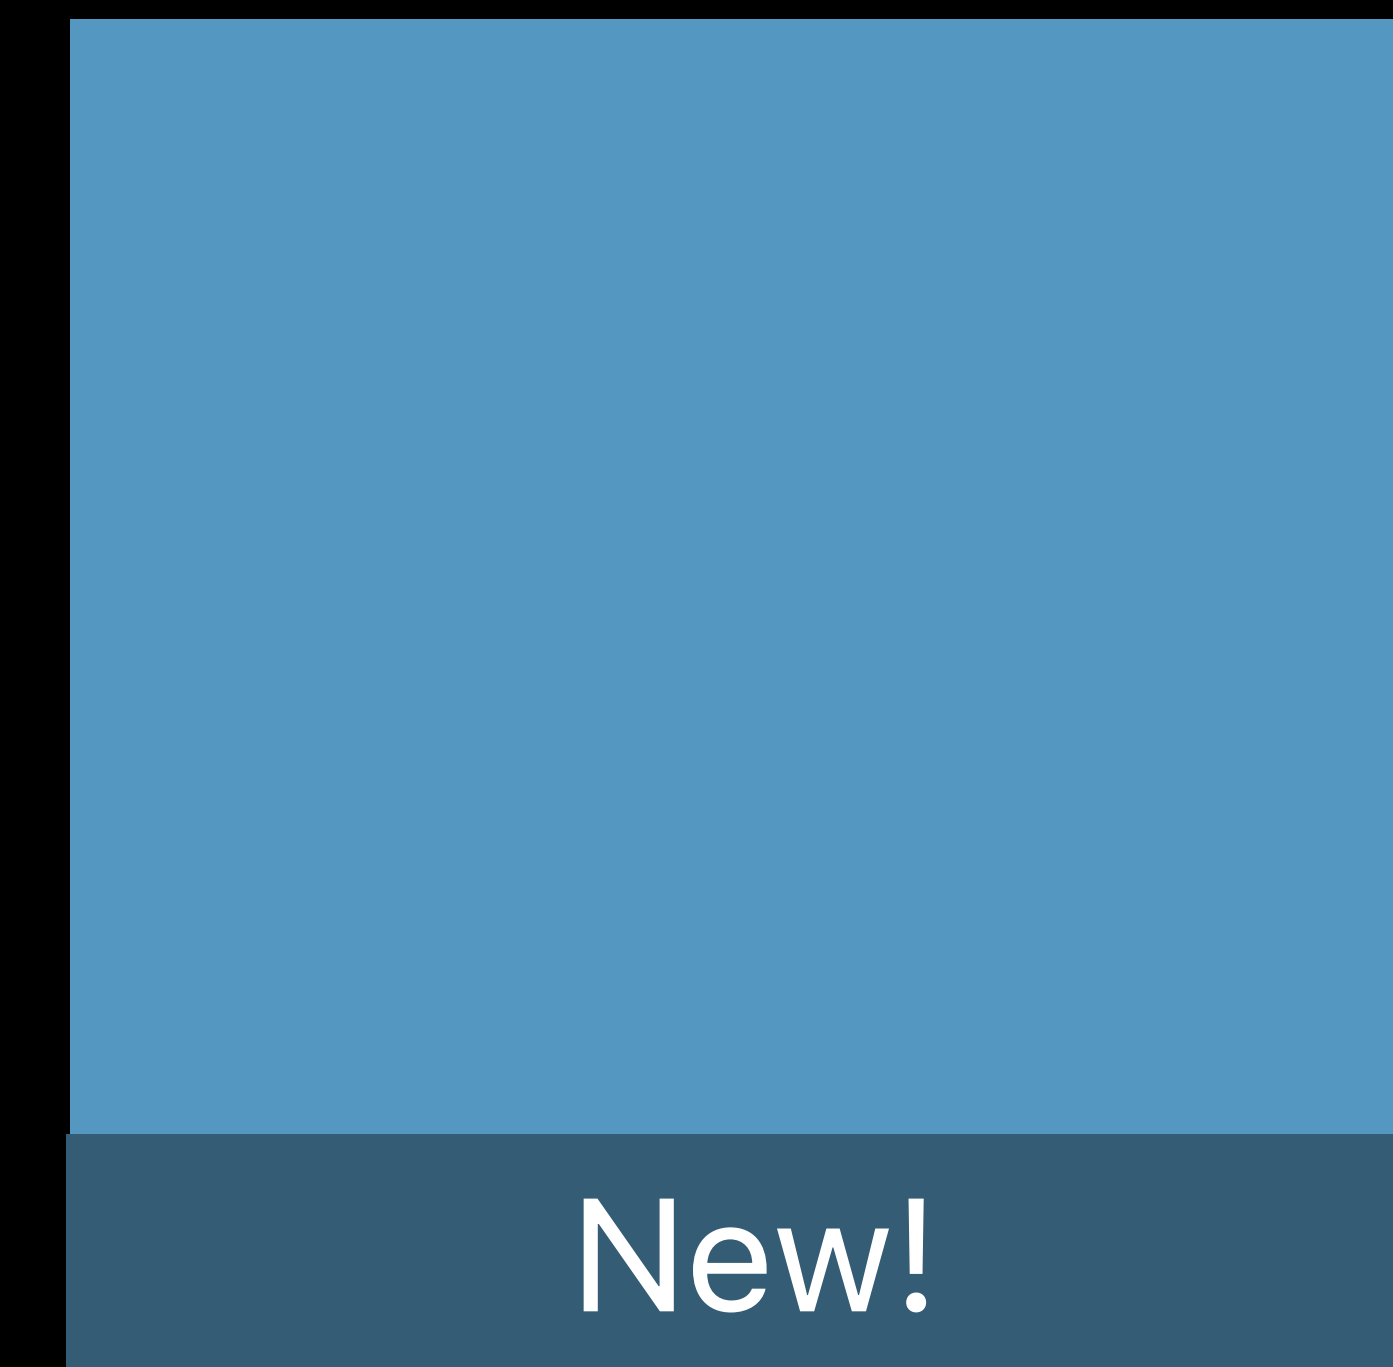
Update file system node (4K)

Update object map (4K)

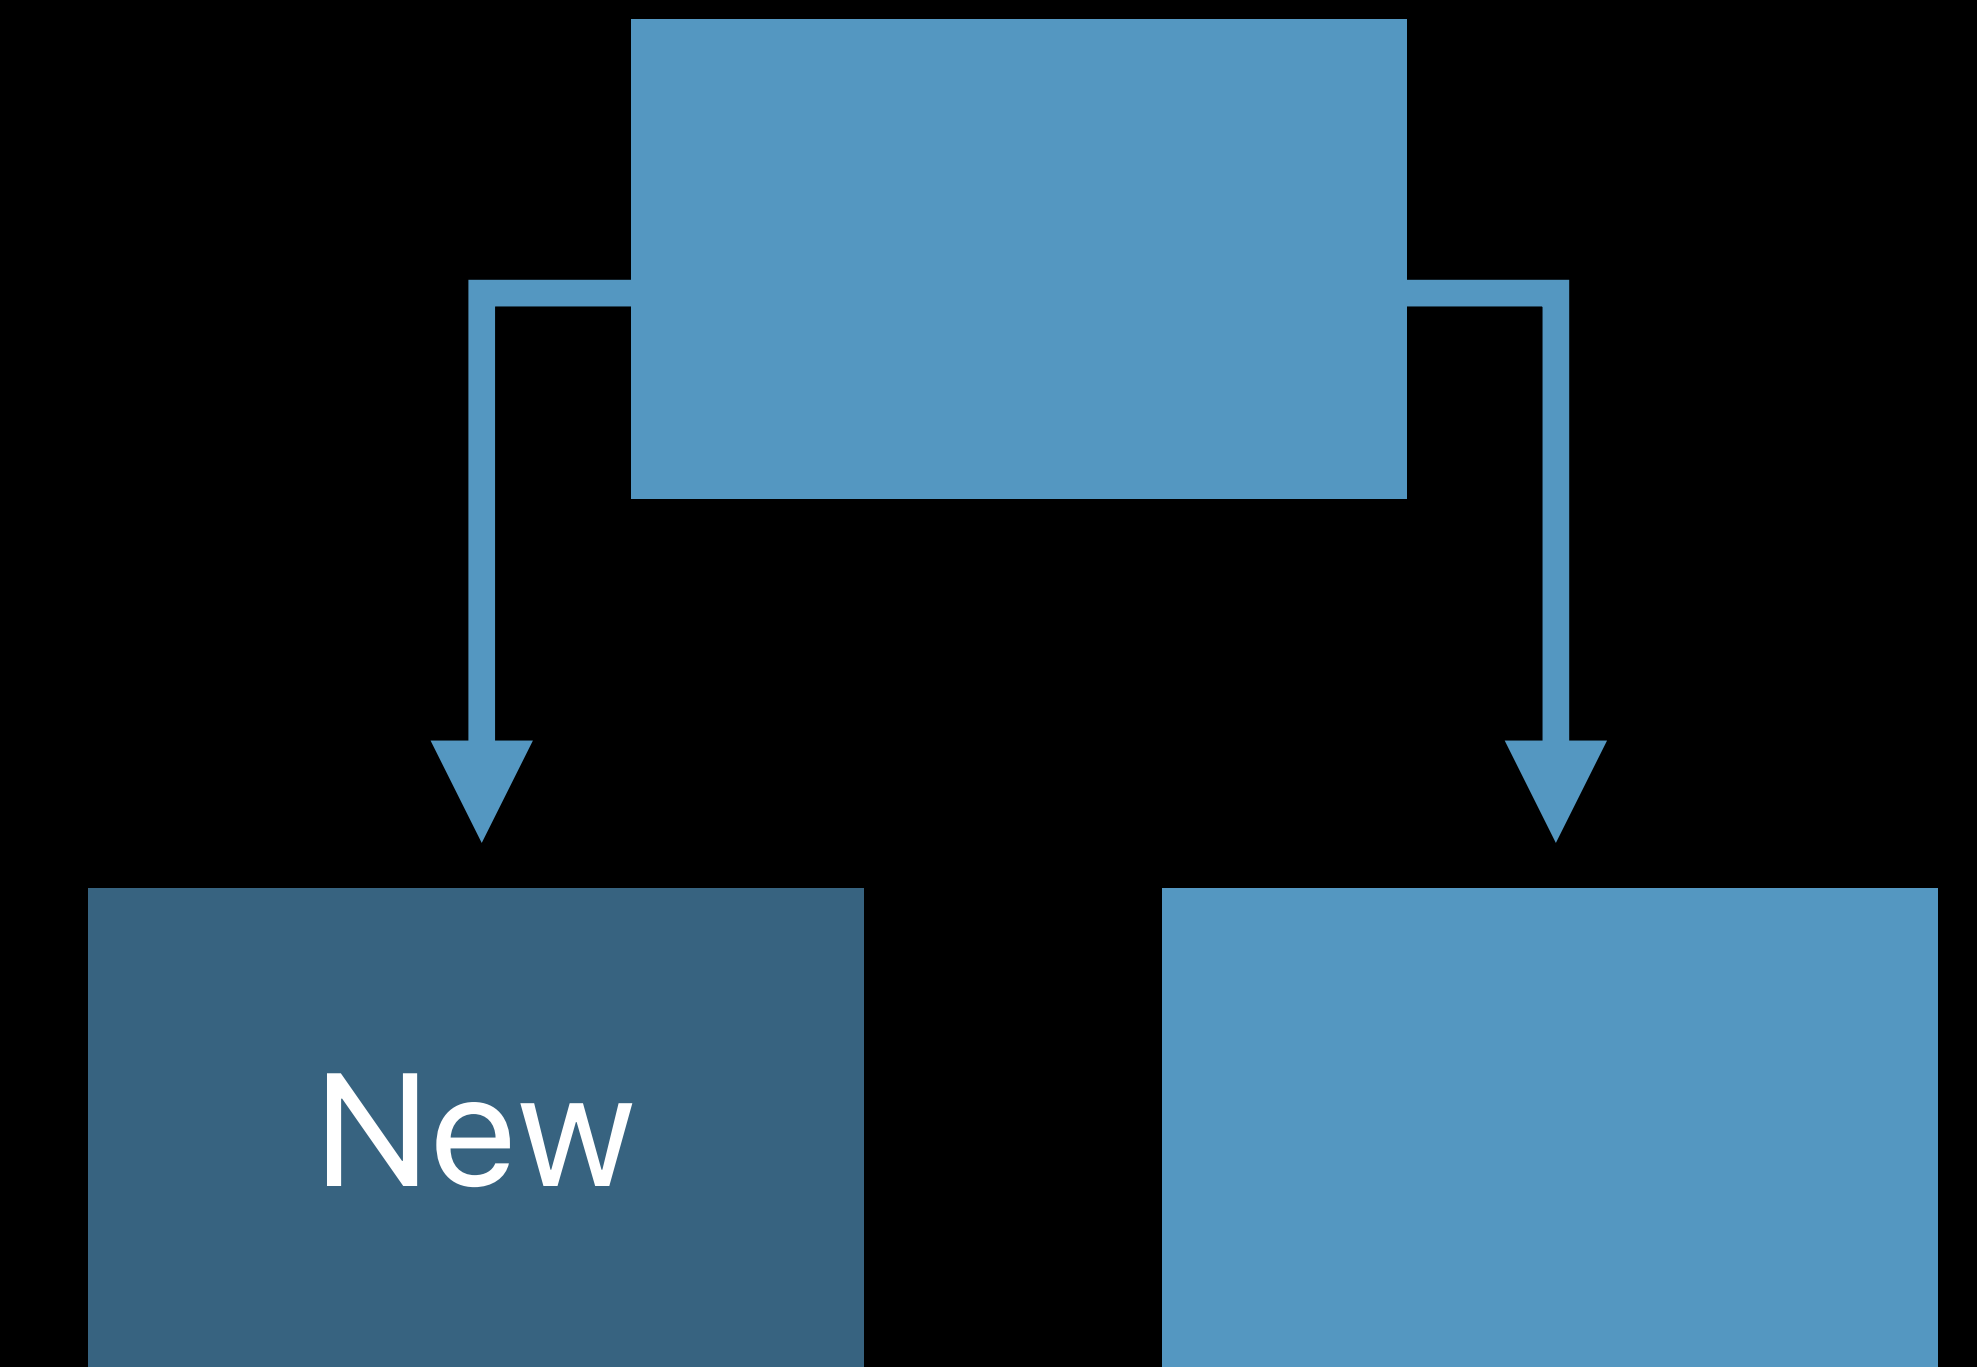
Total Metadata: 8K

File itself

Object Map



File System Tree



4K

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST



# File System Metadata

Writing a 240 byte NSDictionary to a file

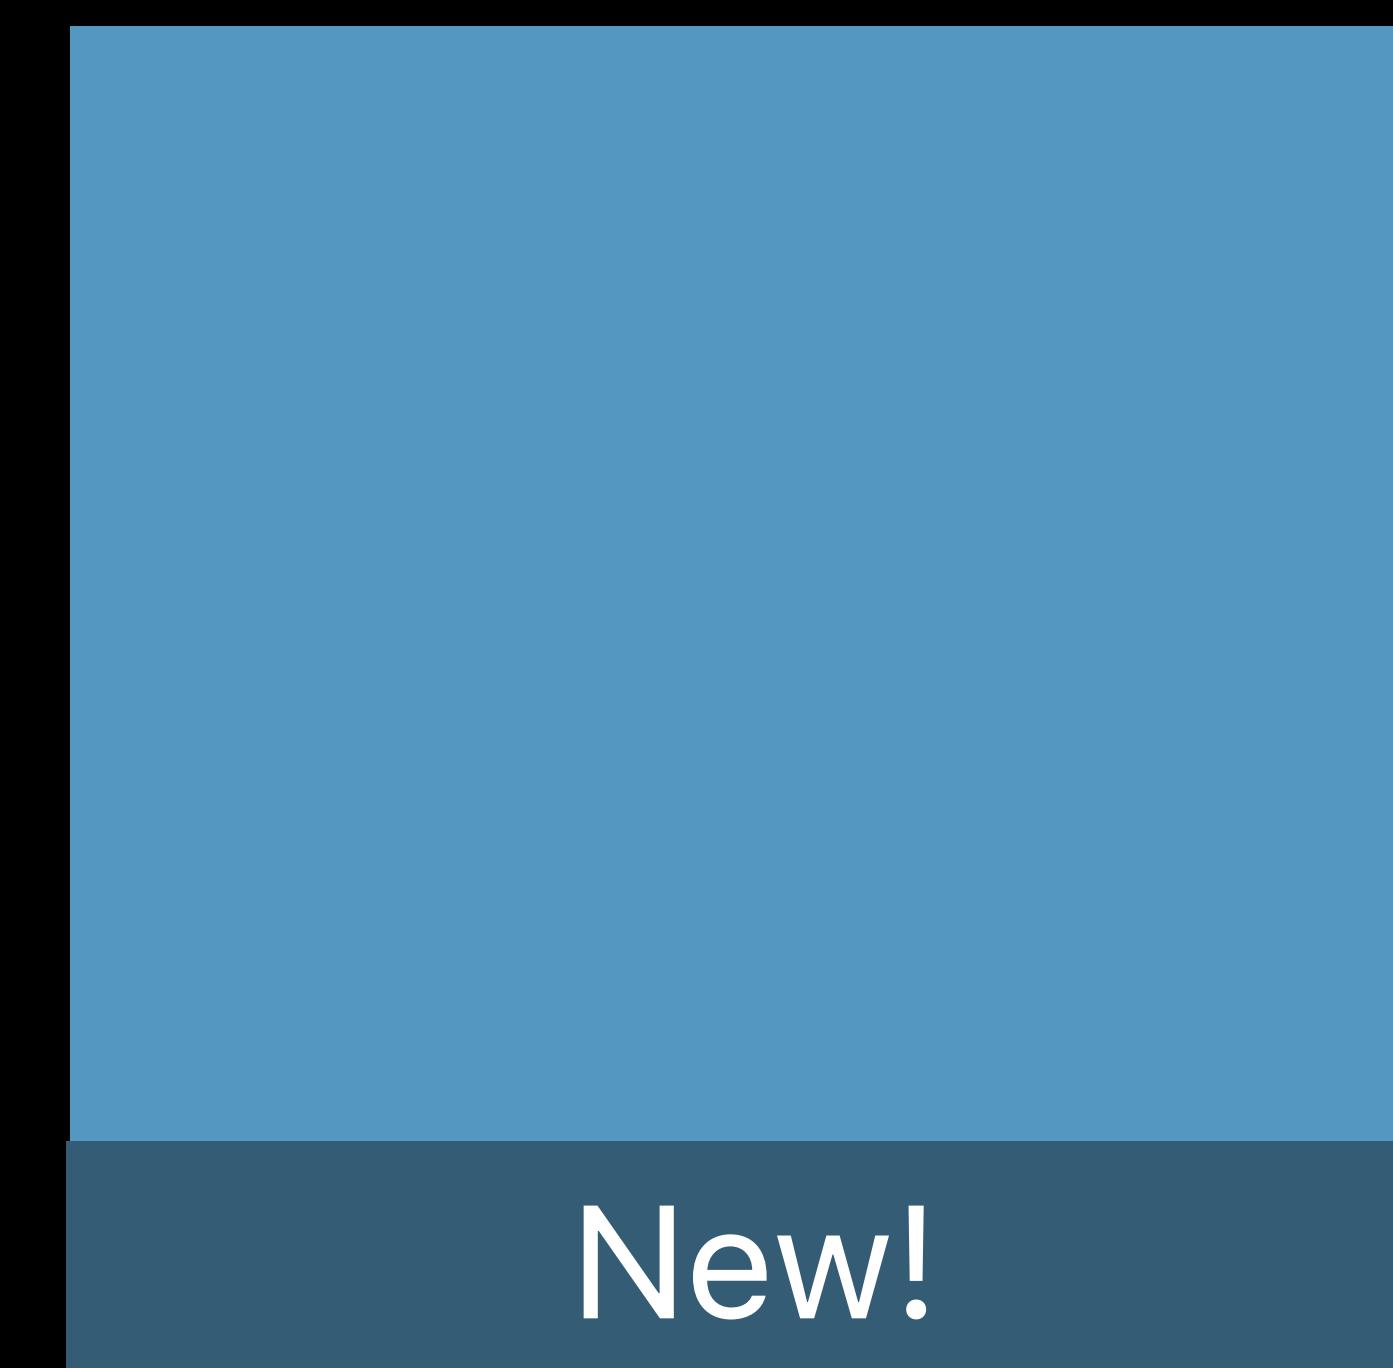
Update file system node (4K)

Update object map (4K)

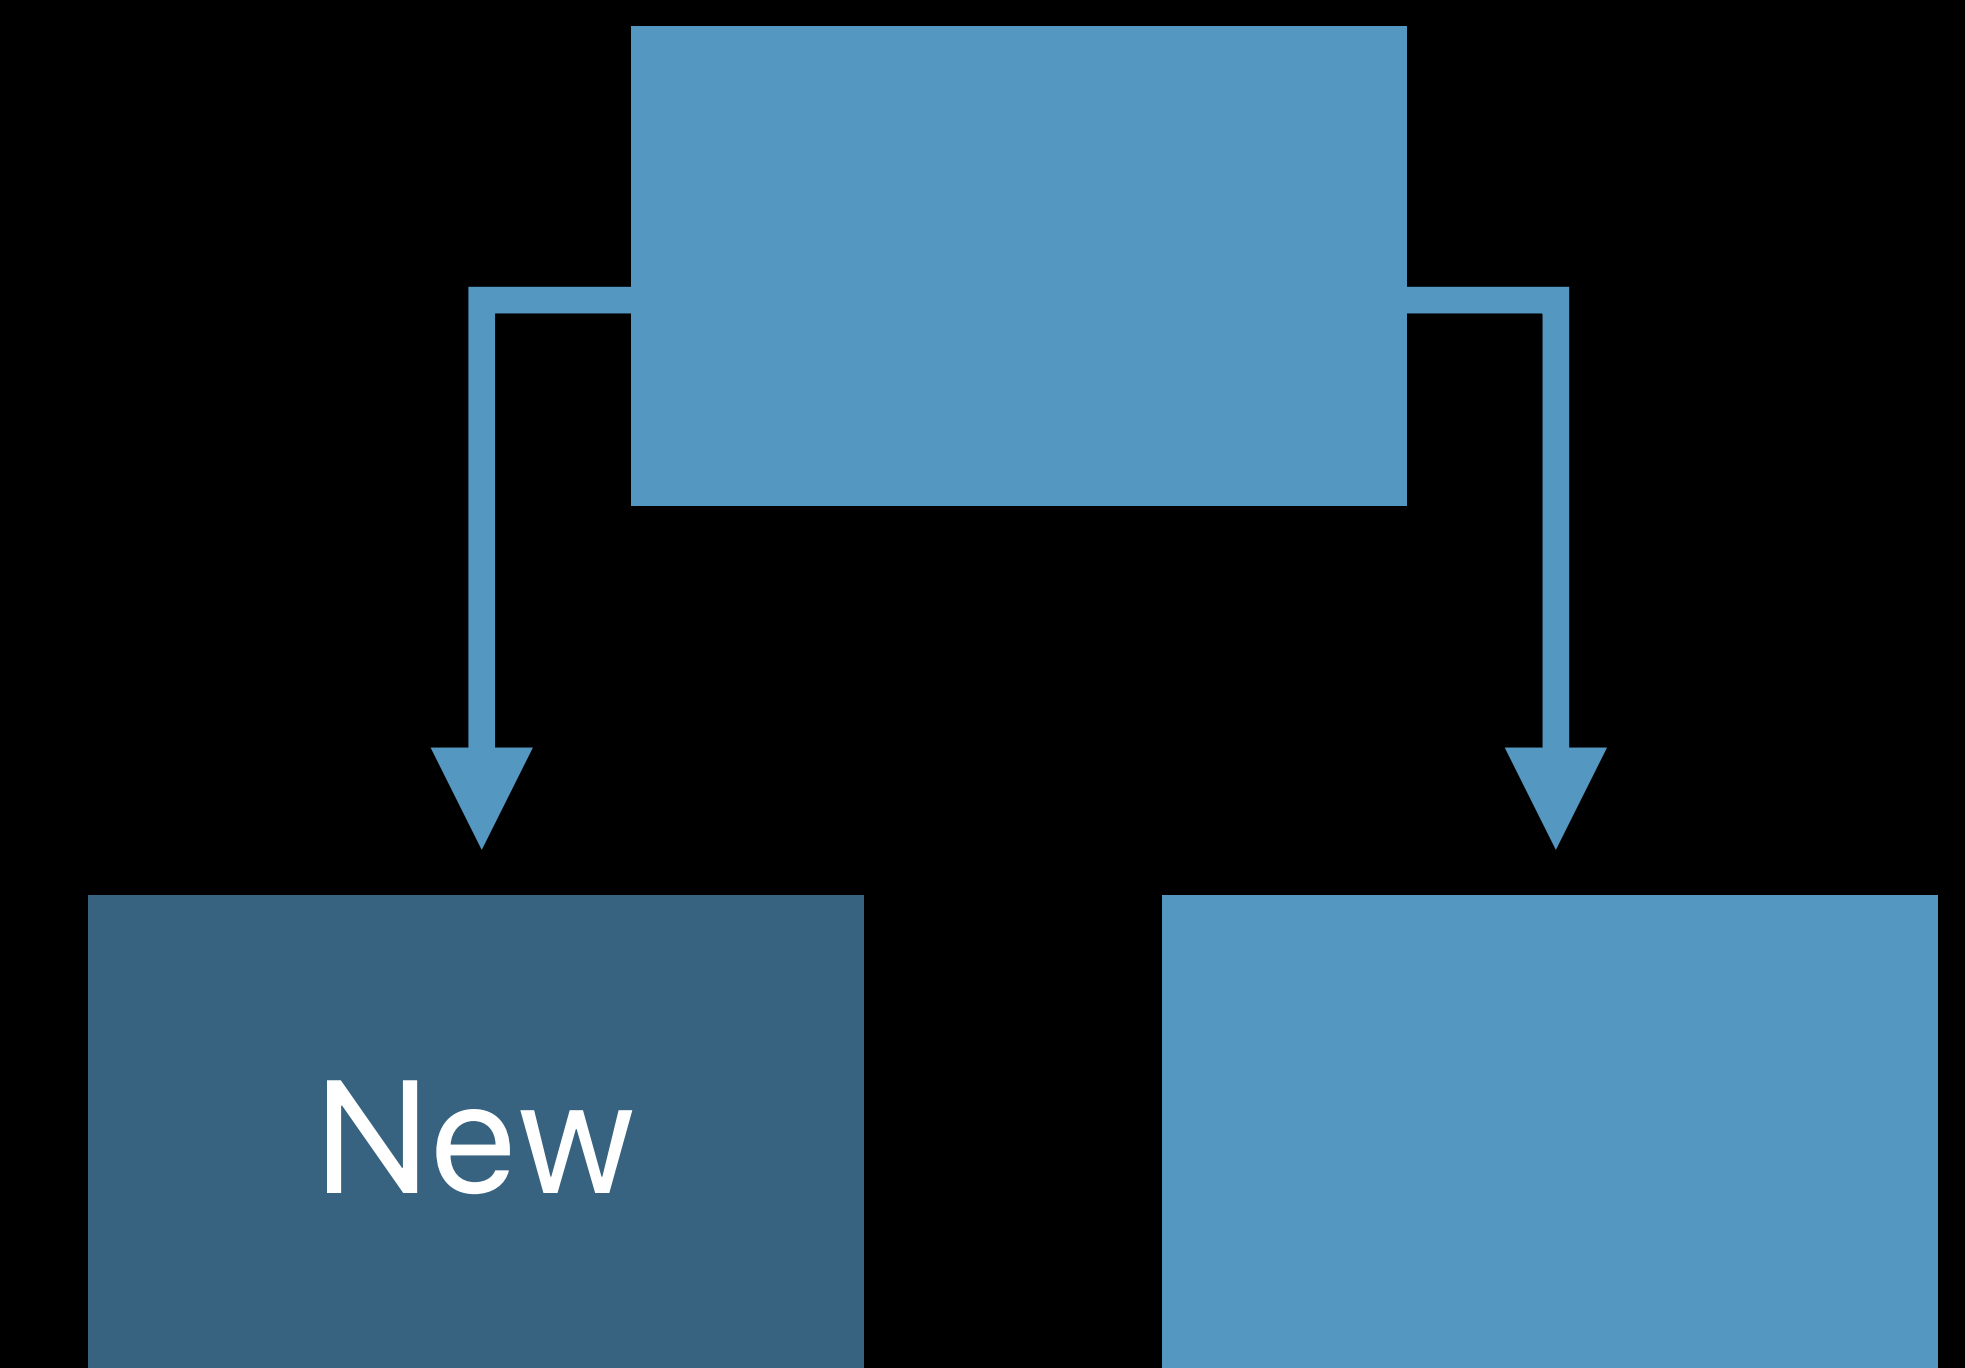
Total Metadata: 8K

File itself (4K)

Object Map



File System Tree



4K

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST

# File System Metadata

Writing a 240 byte NSDictionary to a file

Update file system node (4K)

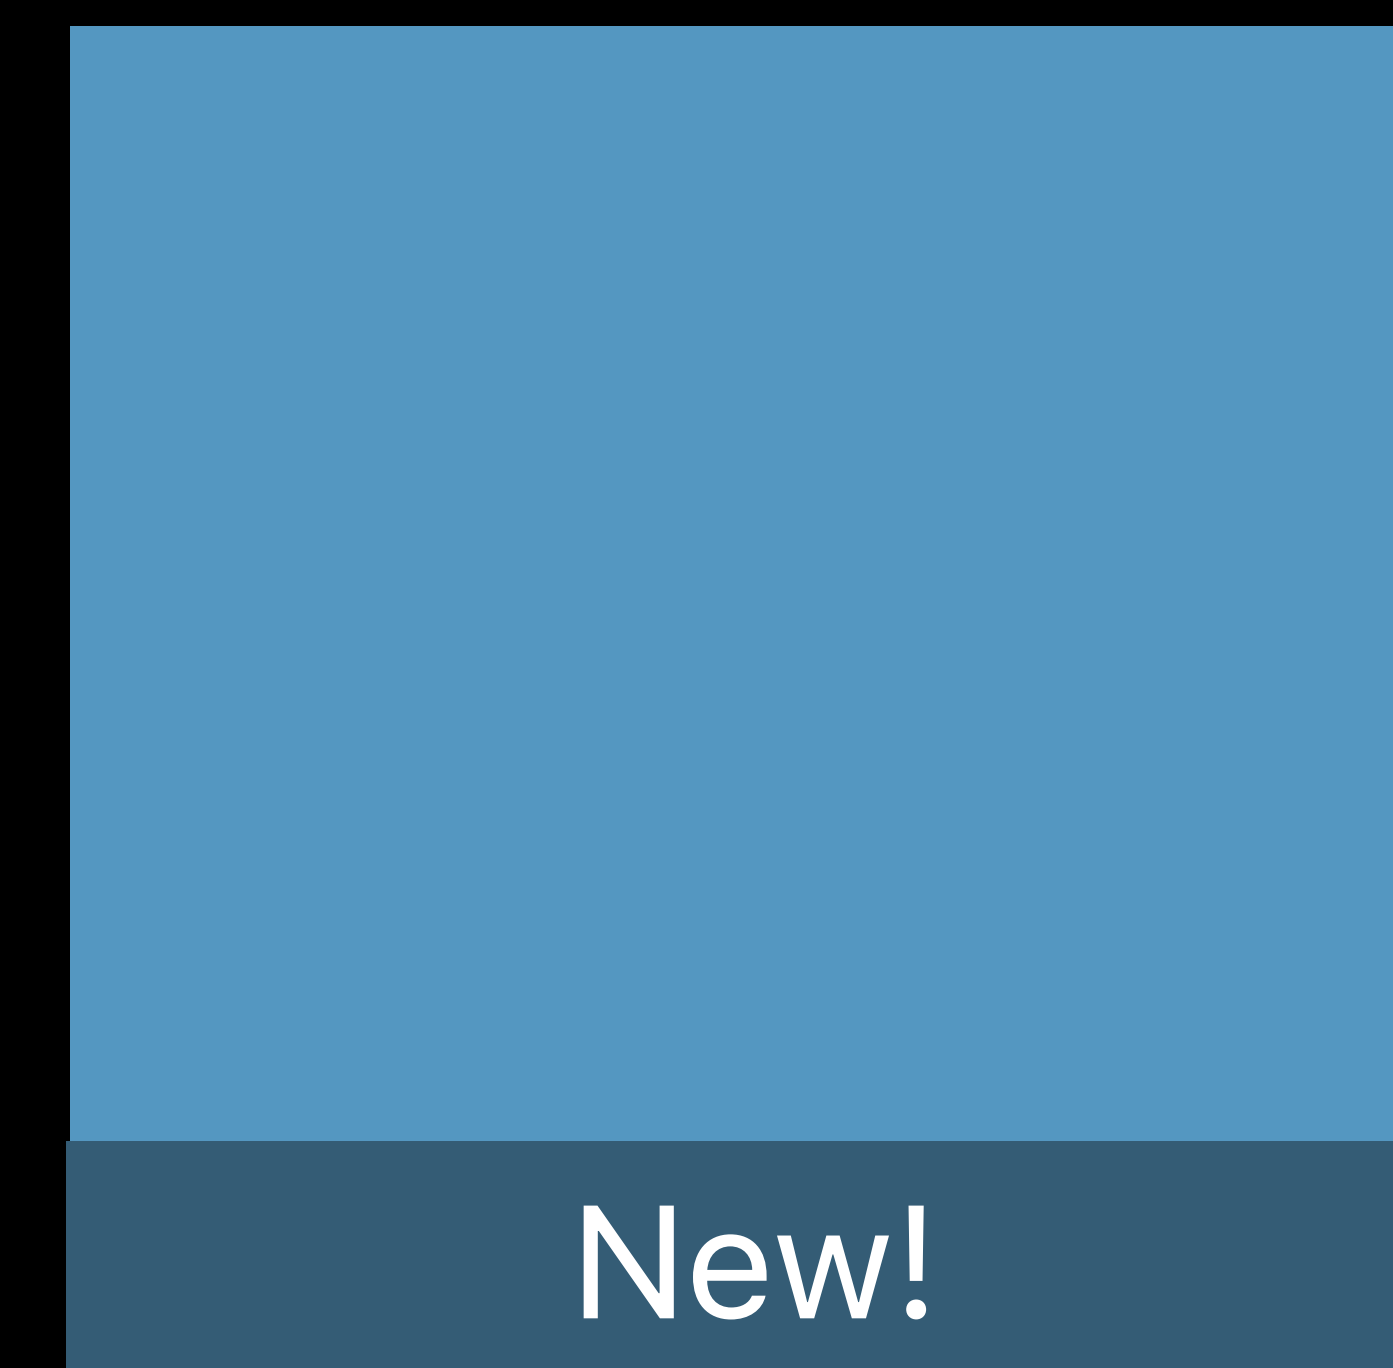
Update object map (4K)

Total Metadata: 8K

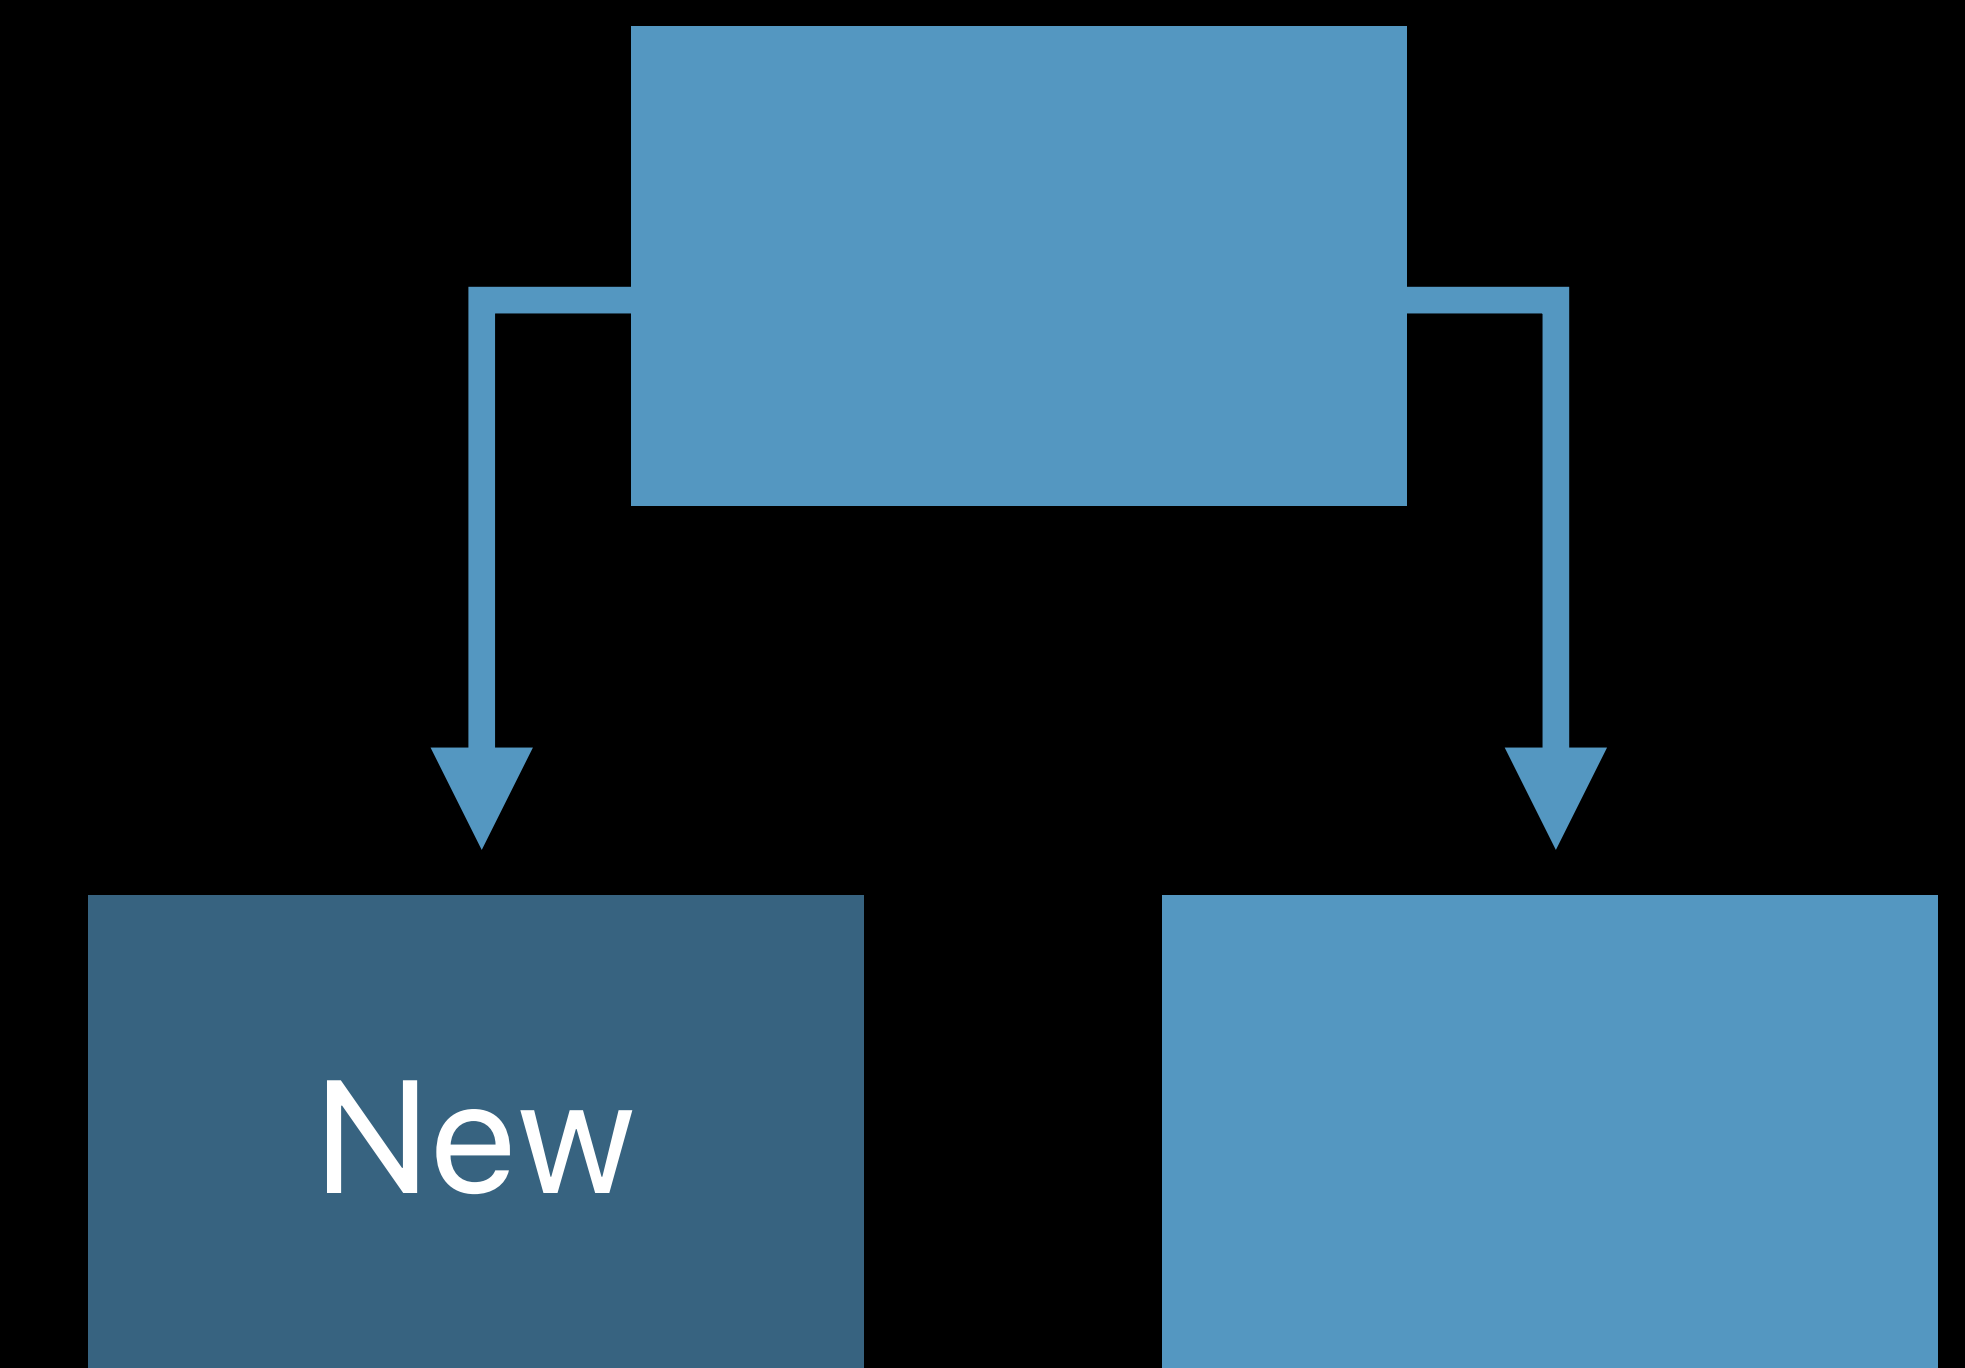
File itself (4K)

Total data: 12K

Object Map



File System Tree



4K

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST



# File System Metadata

Book keeping workload for common operations

Create a file: 8K

# File System Metadata

Book keeping workload for common operations

Create a file: 8K

Delete a file: 8K



# File System Metadata

Book keeping workload for common operations

Create a file: 8K

Delete a file: 8K

Rename a file: 16K

# File System Metadata

Book keeping workload for common operations

Create a file: 8K

Delete a file: 8K

Rename a file: 16K

Modifying a file: 8K



File system metadata  
updates are not free.

# File System Metadata

Efficient non-persistent files

Create a file



# File System Metadata

Efficient non-persistent files

Create a file

Keep it open and unlinked

# File System Metadata

Efficient non-persistent files

Create a file

Keep it open and unlinked

Do not call `fsync()`



# File System Metadata

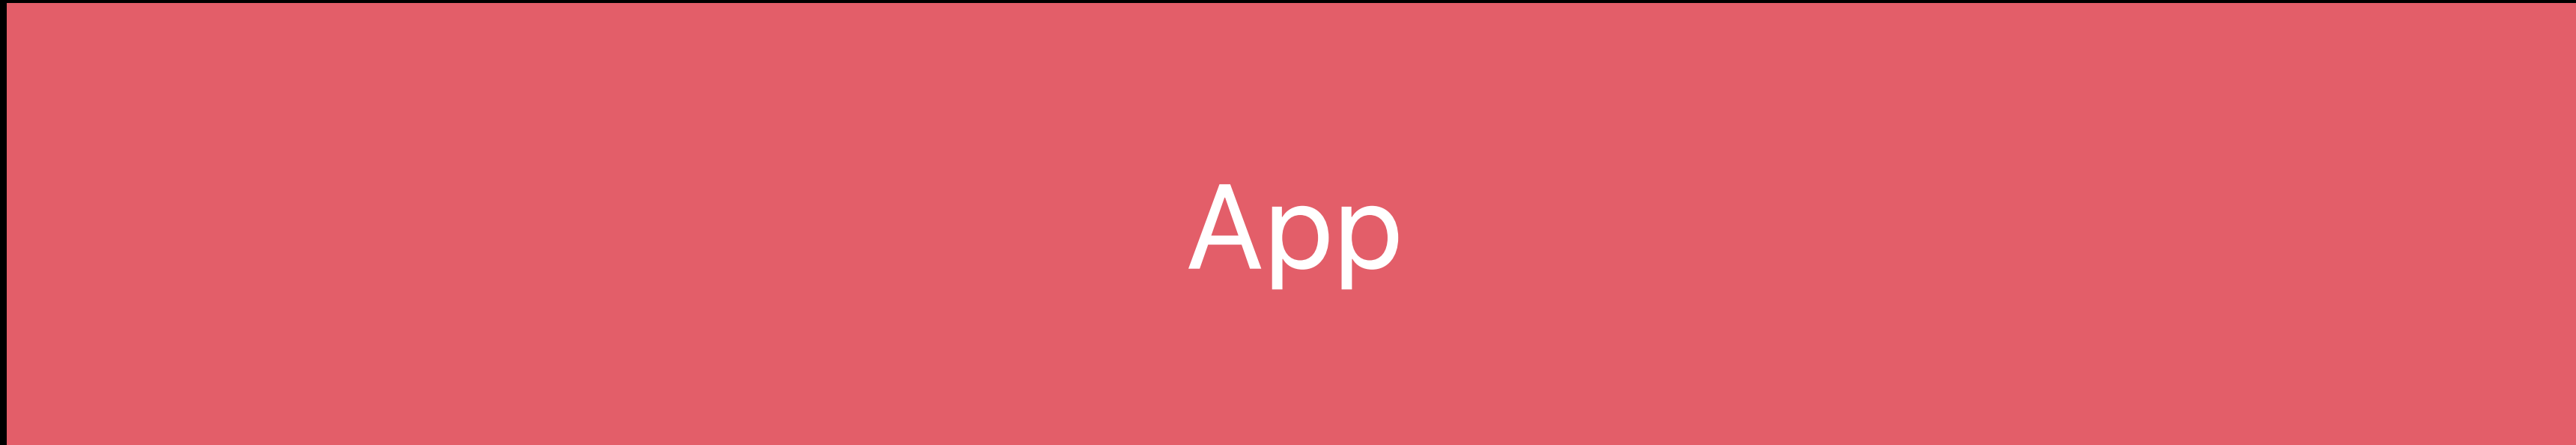
More information on APFS

[developer.apple.com/support/apple-file-system/Apple-File-System-Reference.pdf](https://developer.apple.com/support/apple-file-system/Apple-File-System-Reference.pdf)

**Syncing to Disk**



User Space



App



Kernel



OS Cache



Hardware



Disk Cache



Permanent Storage

User Space



Kernel



Hardware





User Space



Kernel



Logical I/O



Hardware



User Space



Kernel



Logical I/O

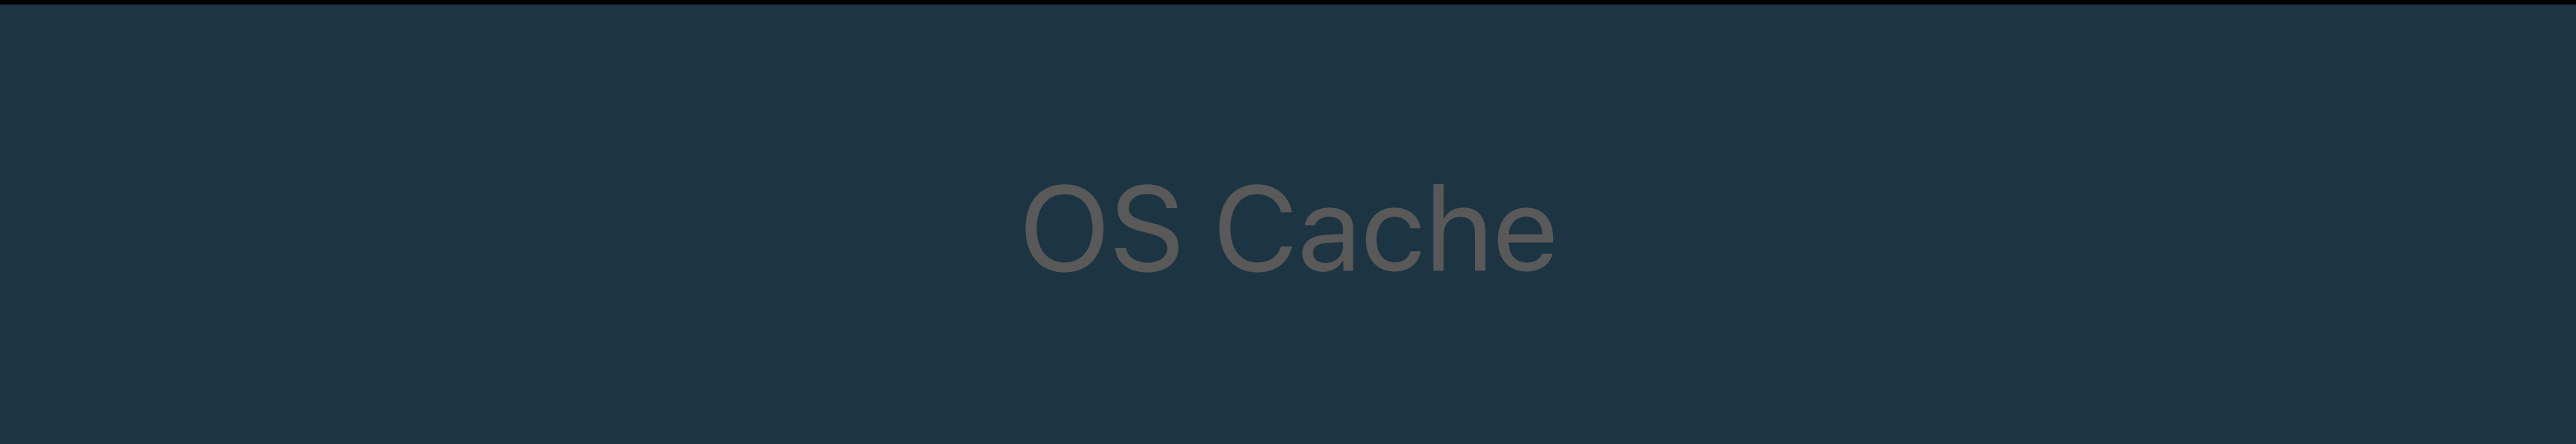
Hardware



User Space



Kernel



Logical I/O



Hardware





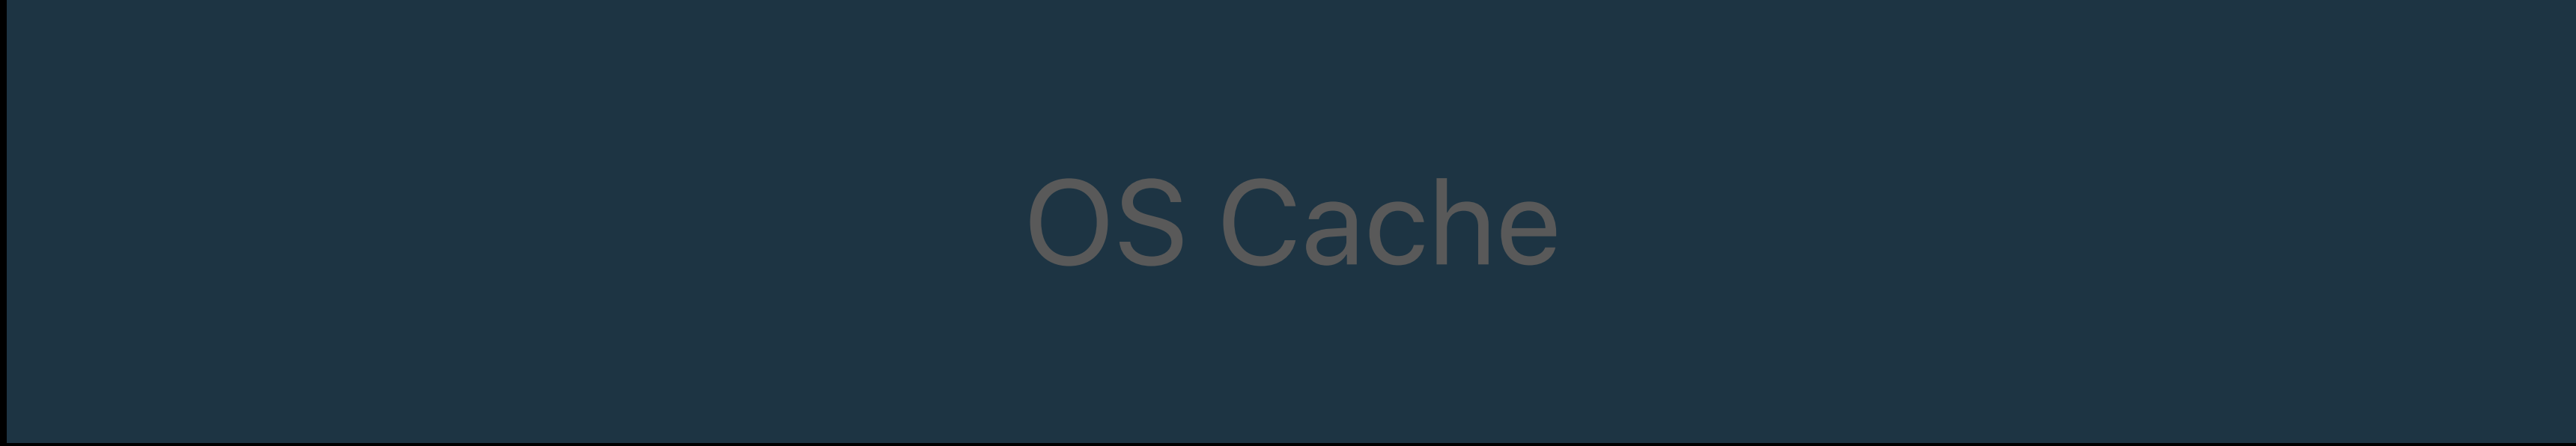
User Space



App



Kernel



OS Cache

Logical I/O



Hardware



Disk Cache

Physical I/O



Permanent Storage

User Space



Kernel



Logical I/O



Hardware

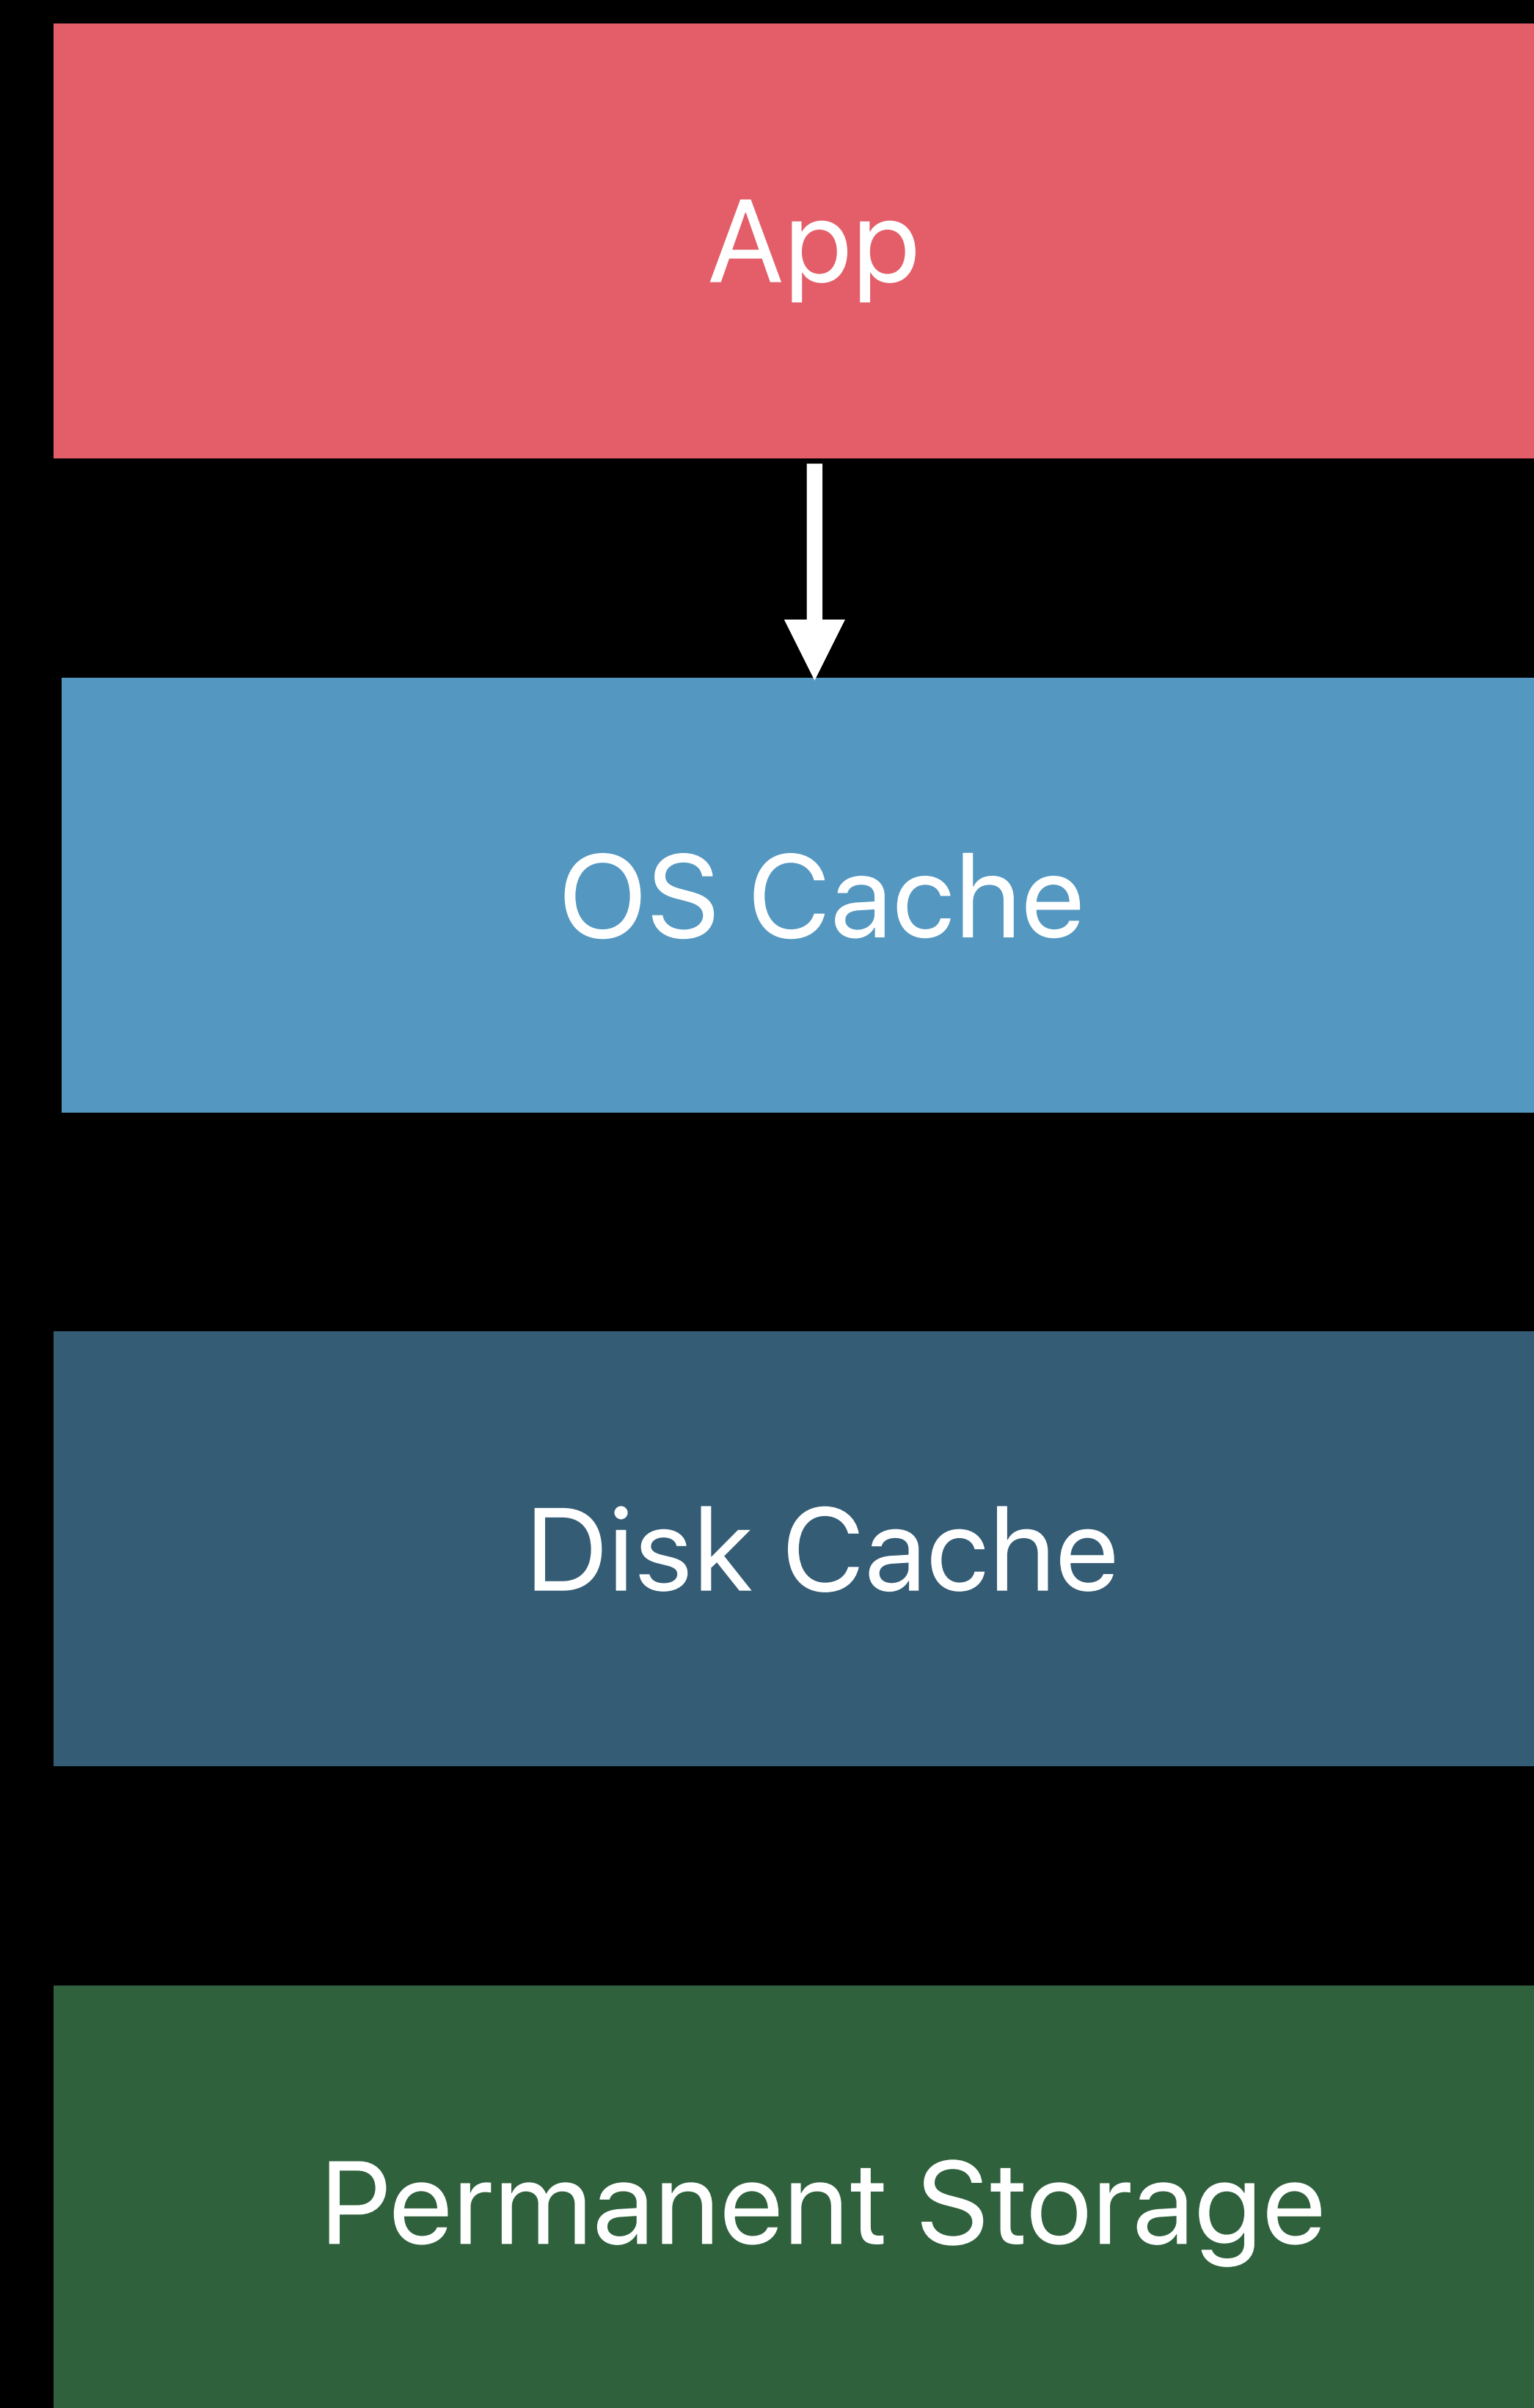


Physical I/O



# fsync()

Flushes cached data to disk

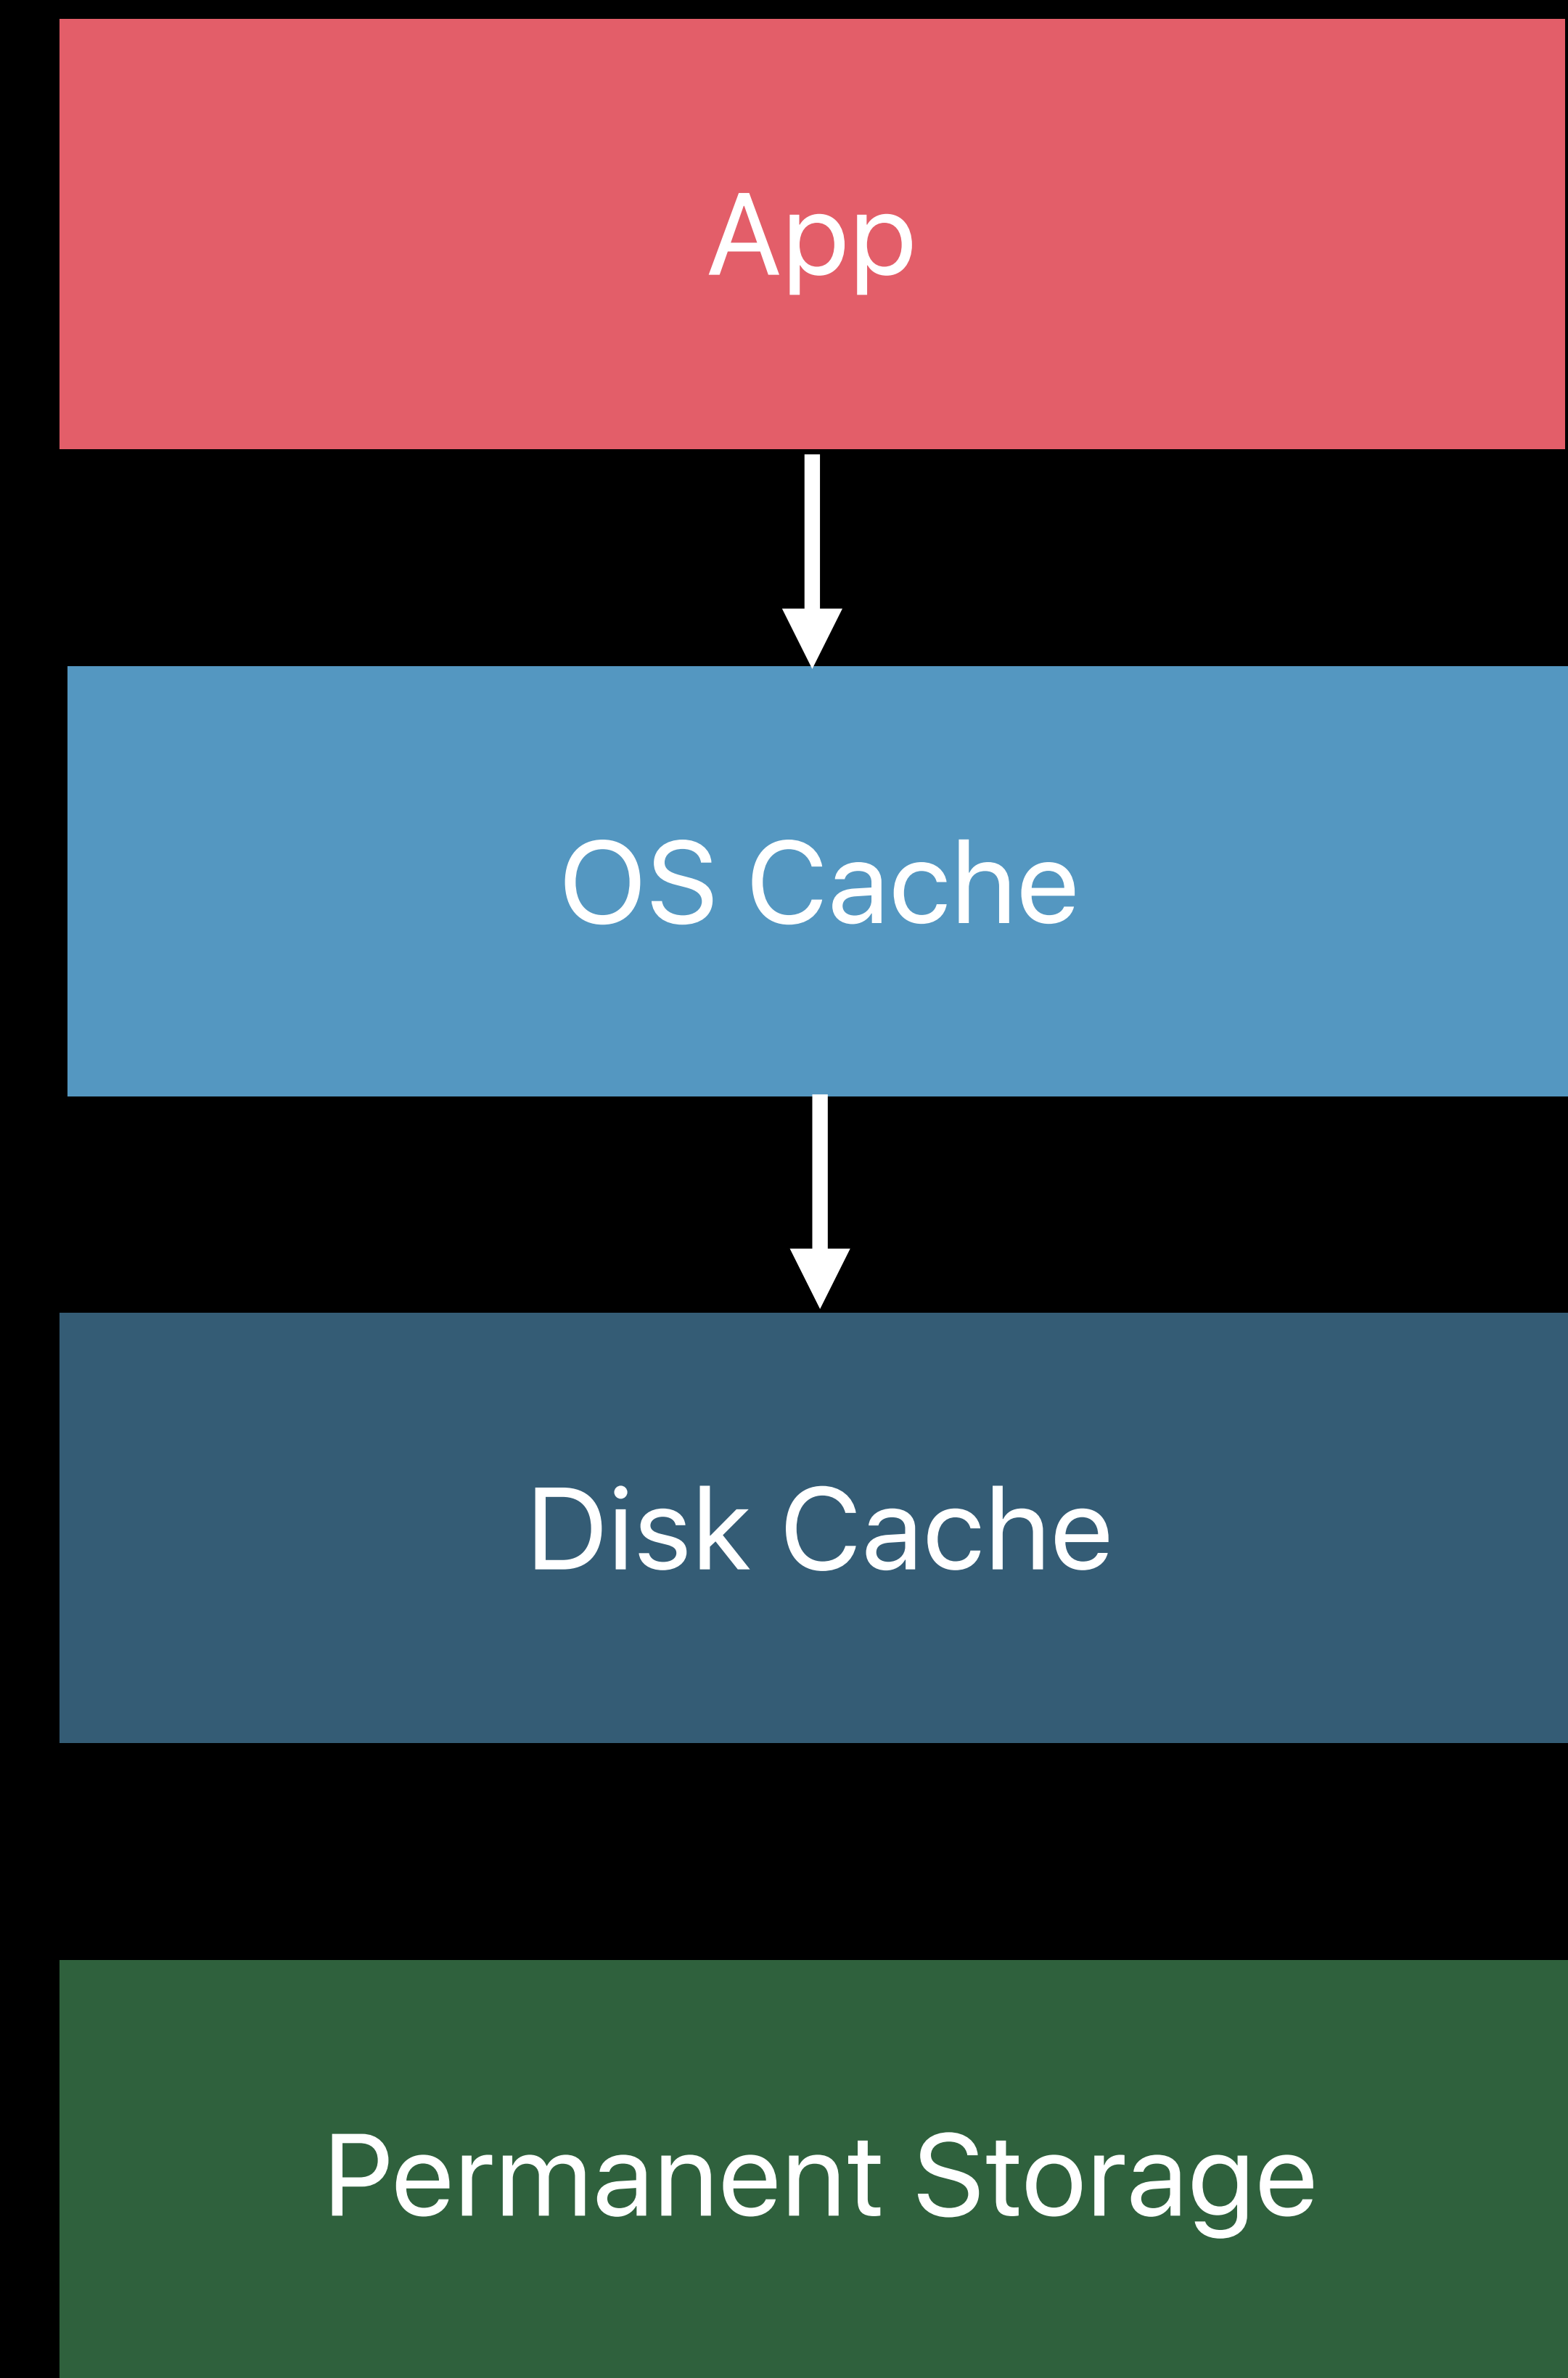




# fsync()

Flushes cached data to disk

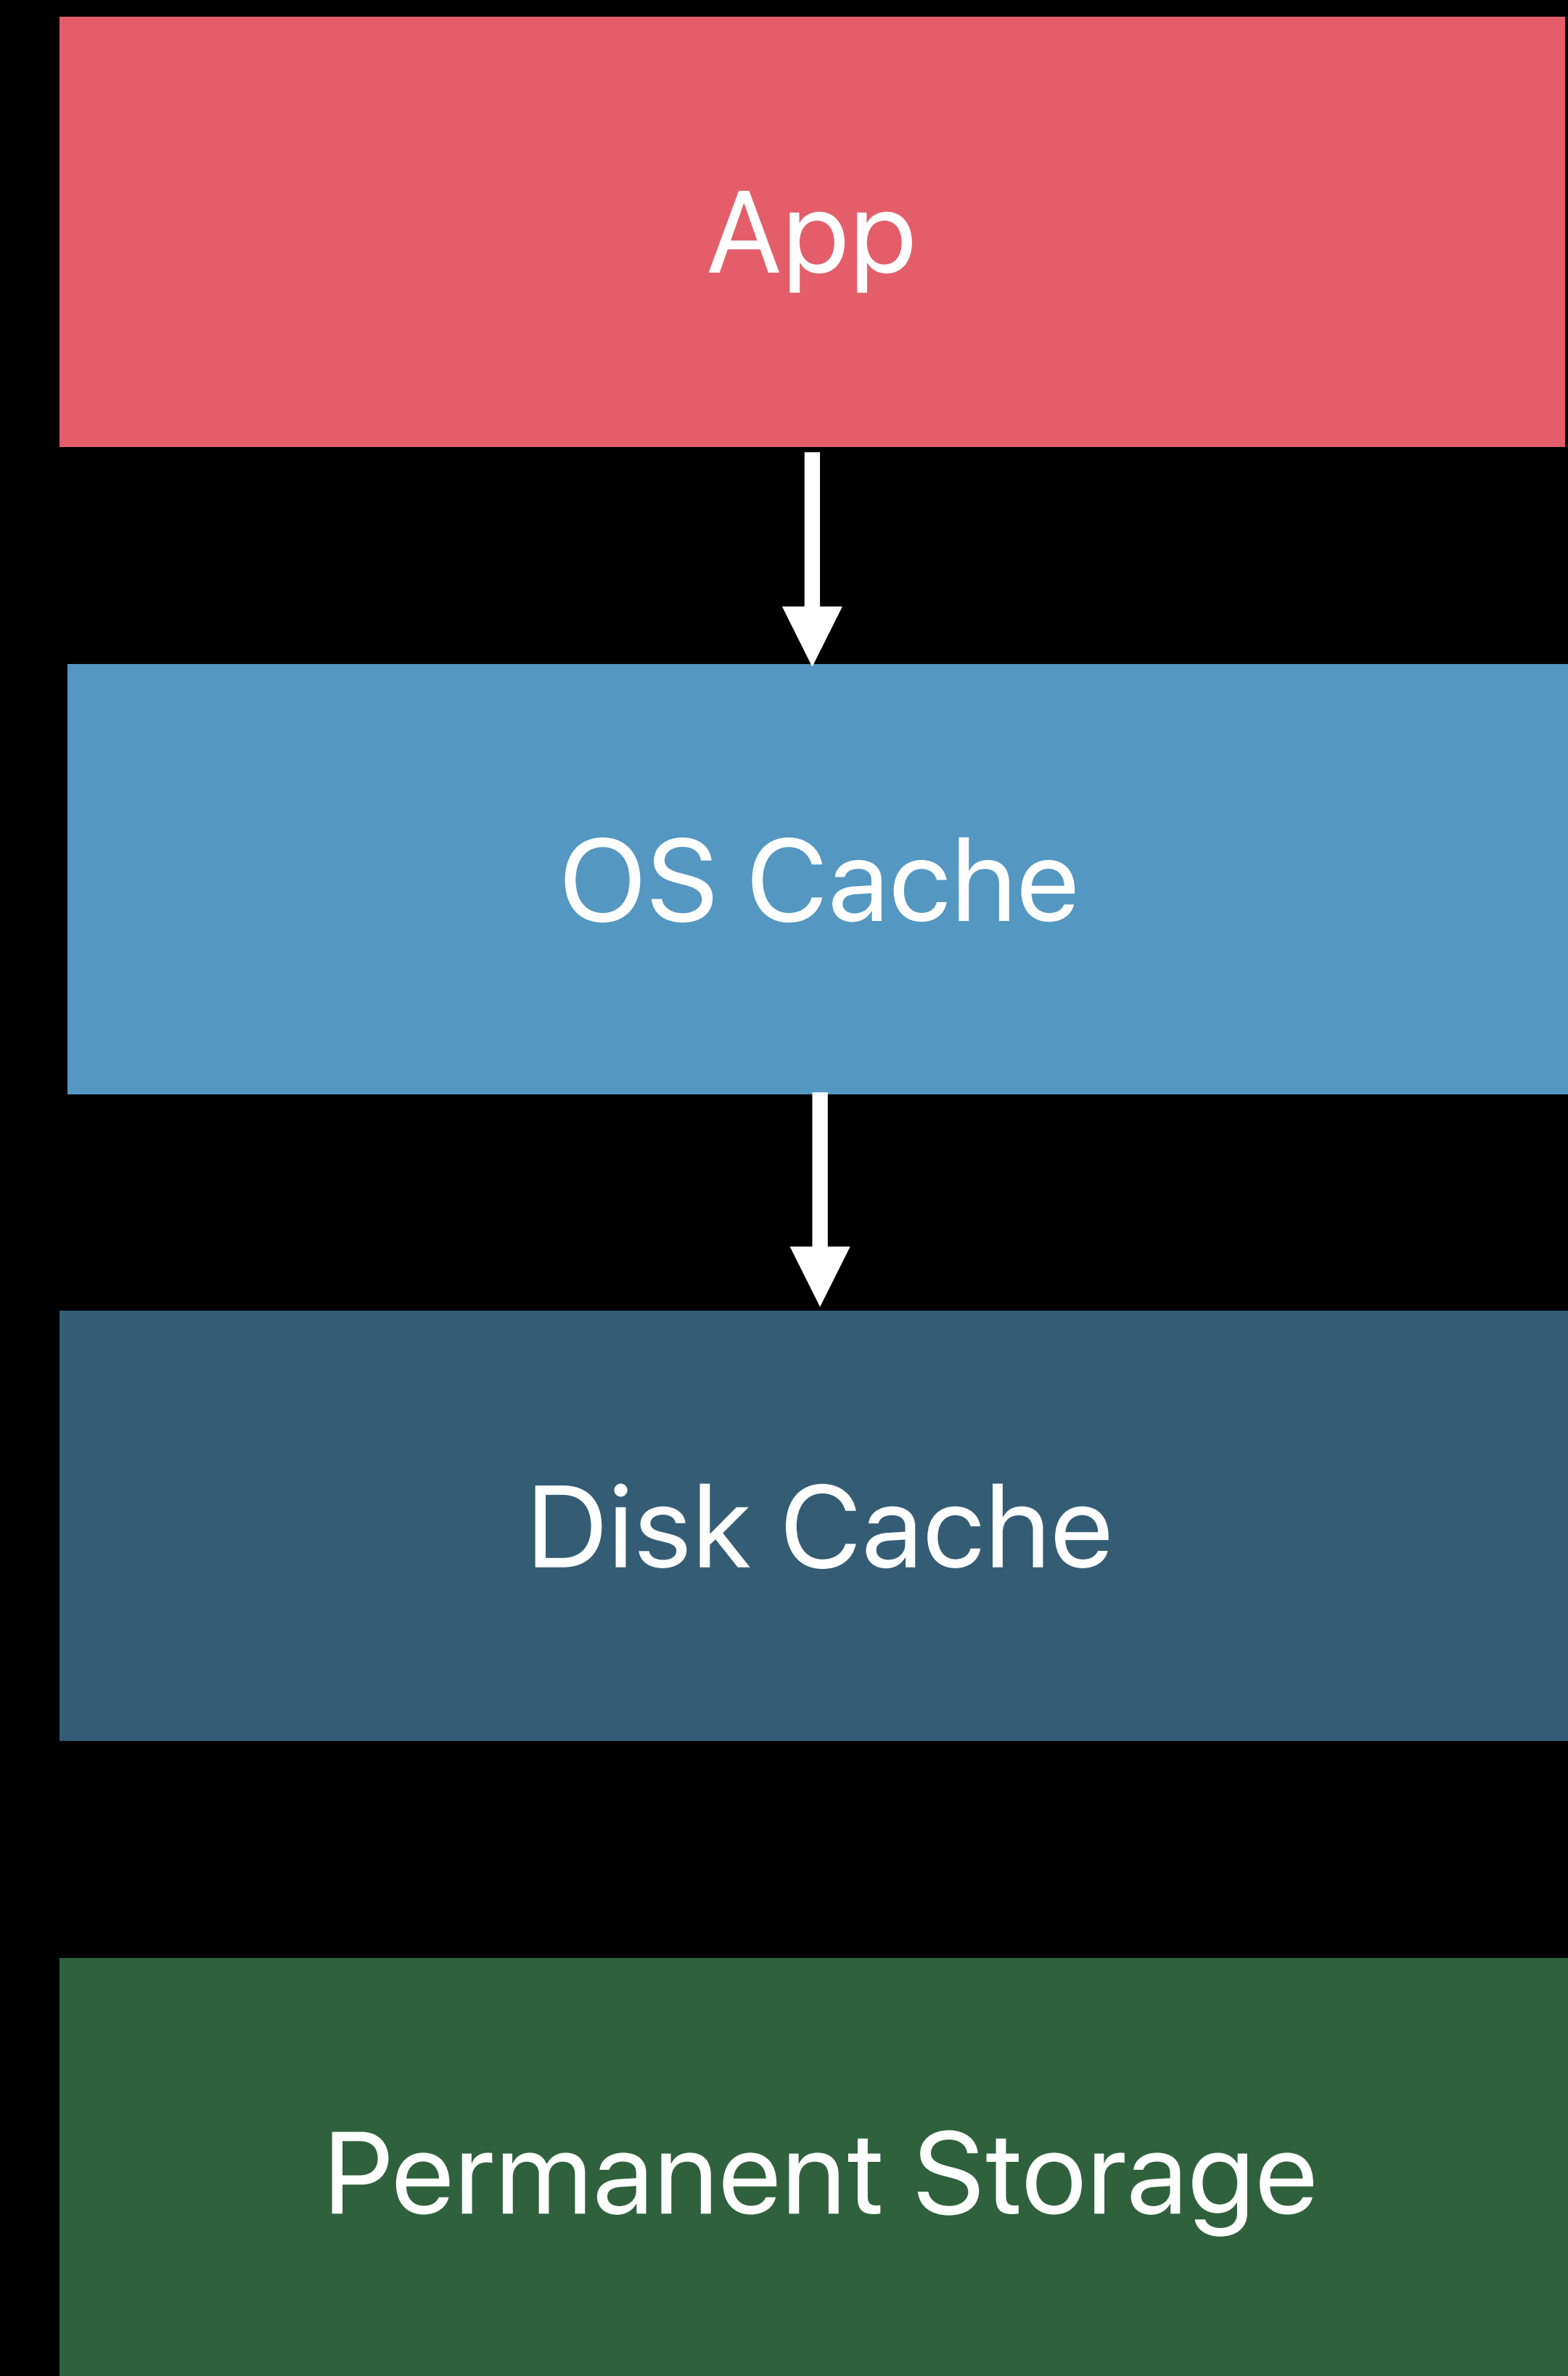
- Moves data from the OS cache to the disk cache



# fsync()

Flushes cached data to disk

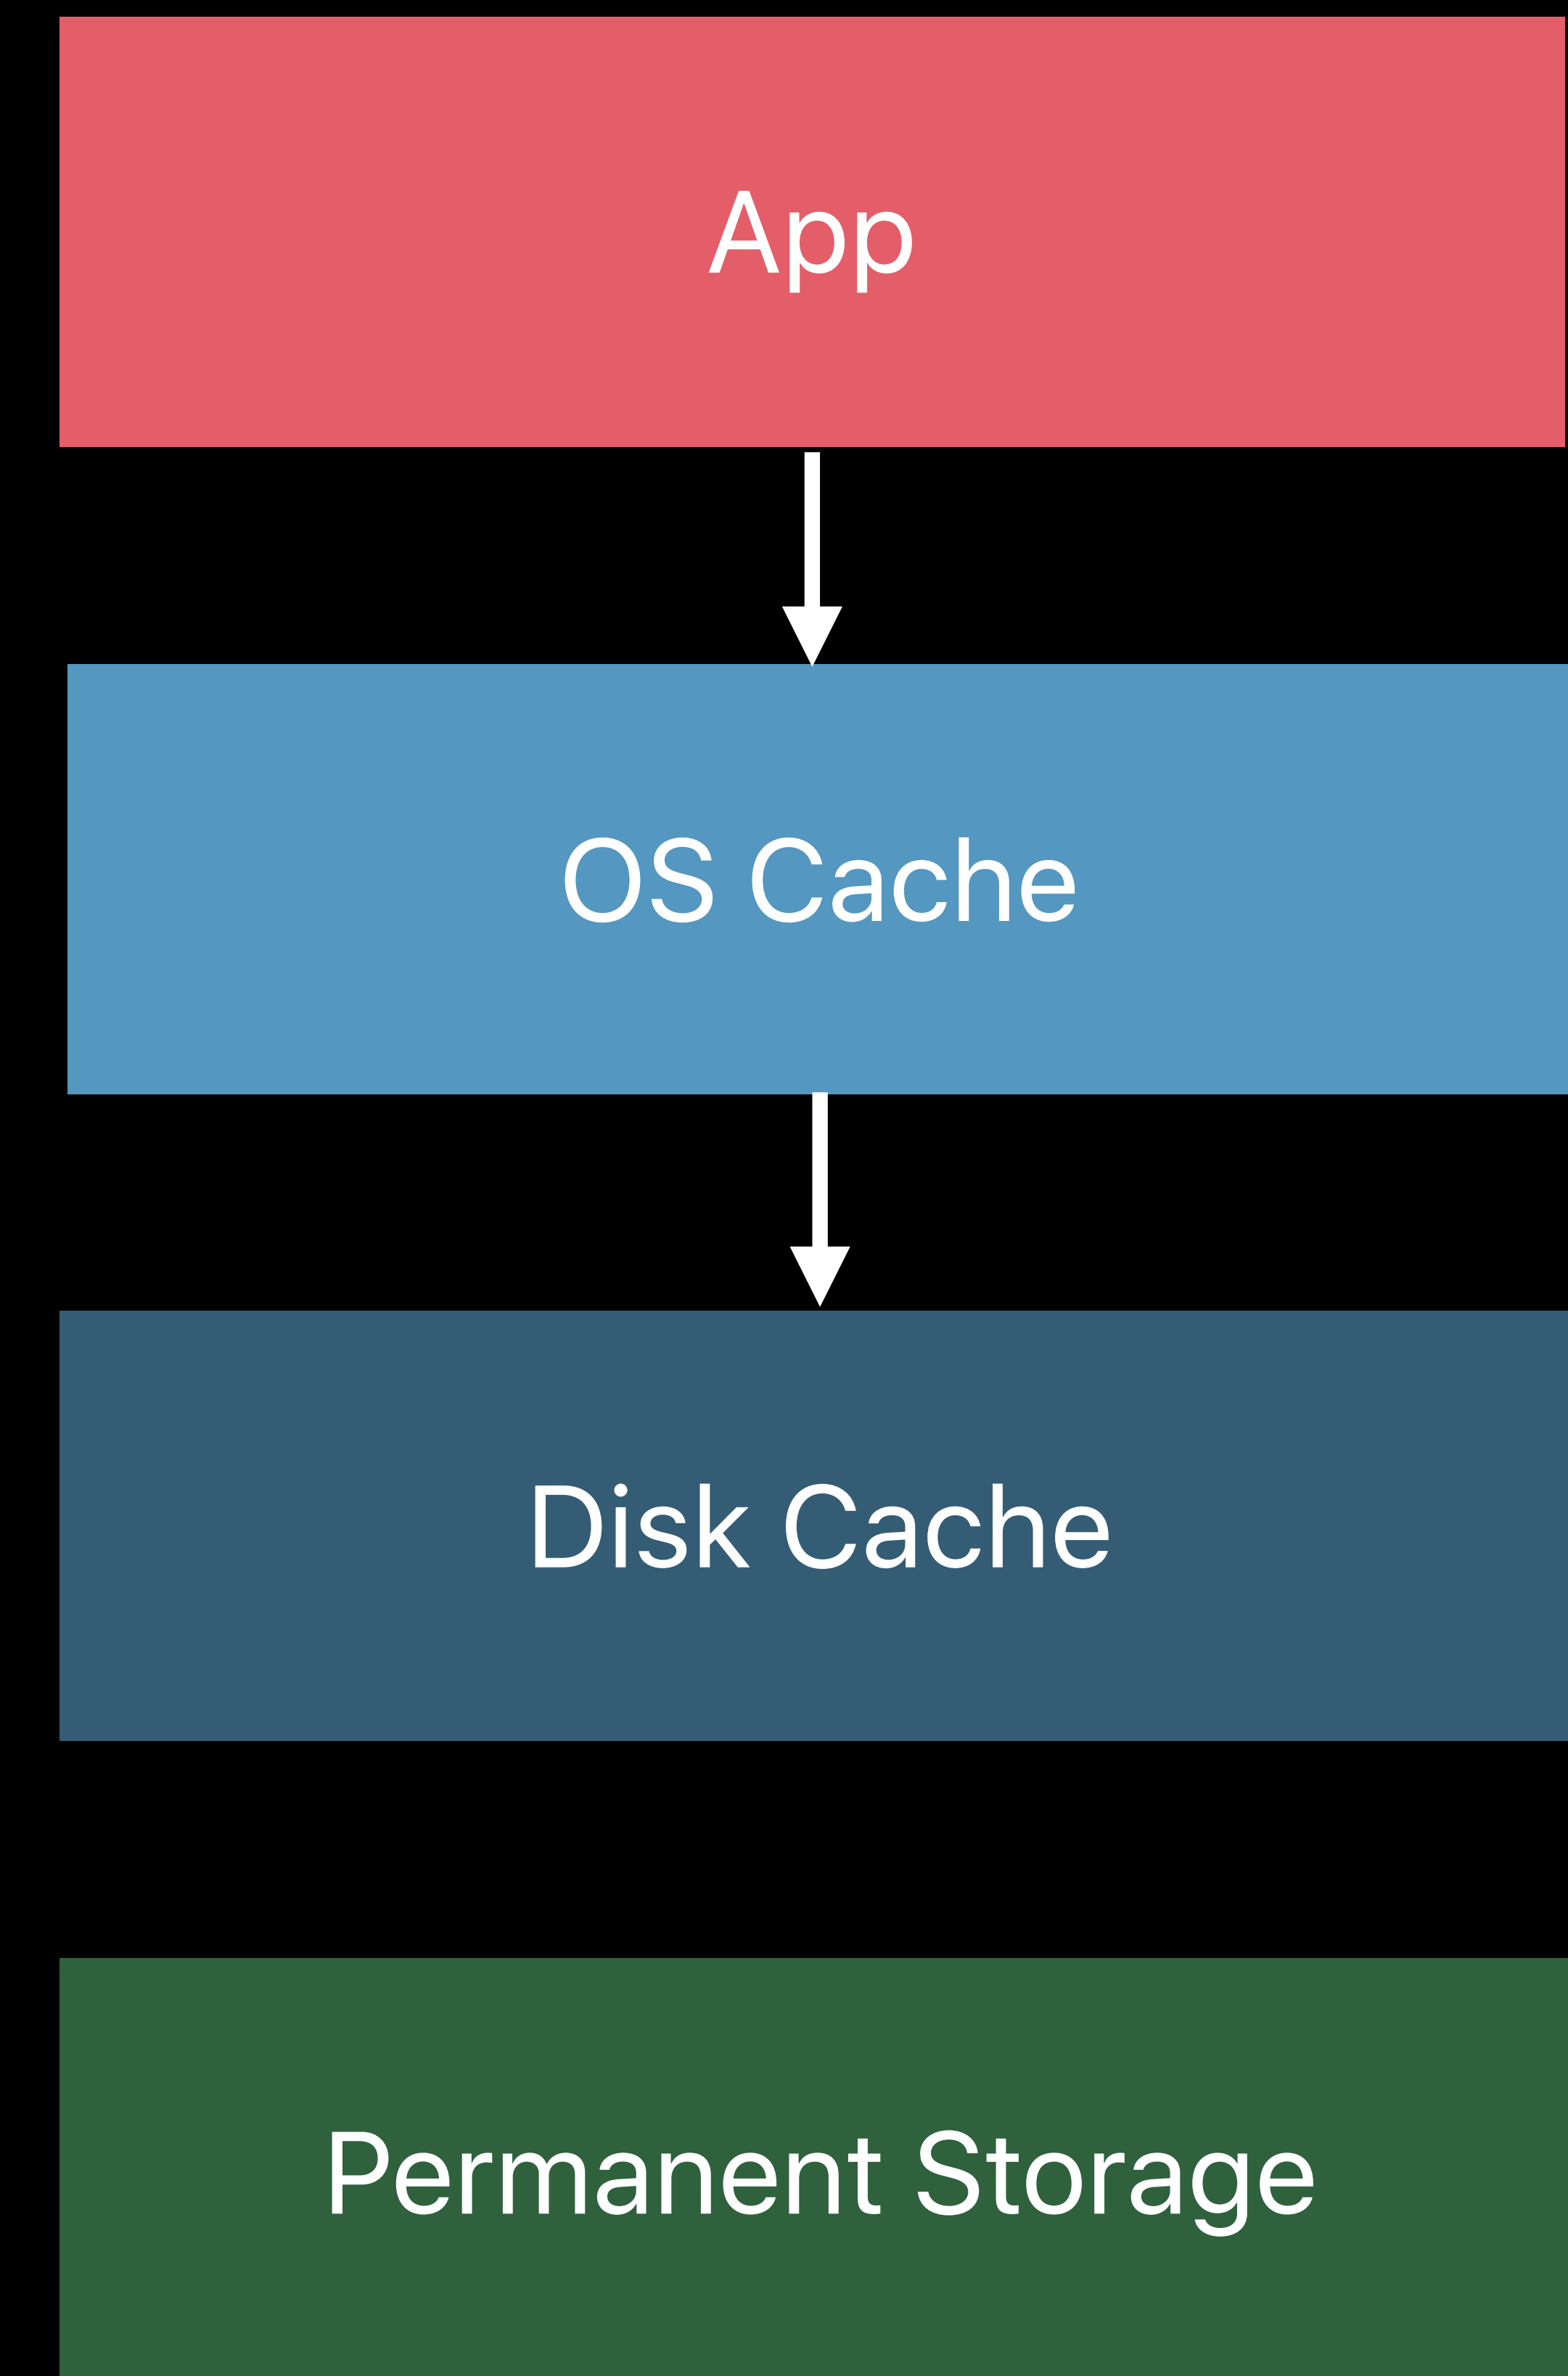
- Moves data from the OS cache to the disk cache
- Data may not be written to permanent storage immediately



# fsync()

Flushes cached data to disk

- Moves data from the OS cache to the disk cache
- Data may not be written to permanent storage immediately
- Does not guarantee write ordering

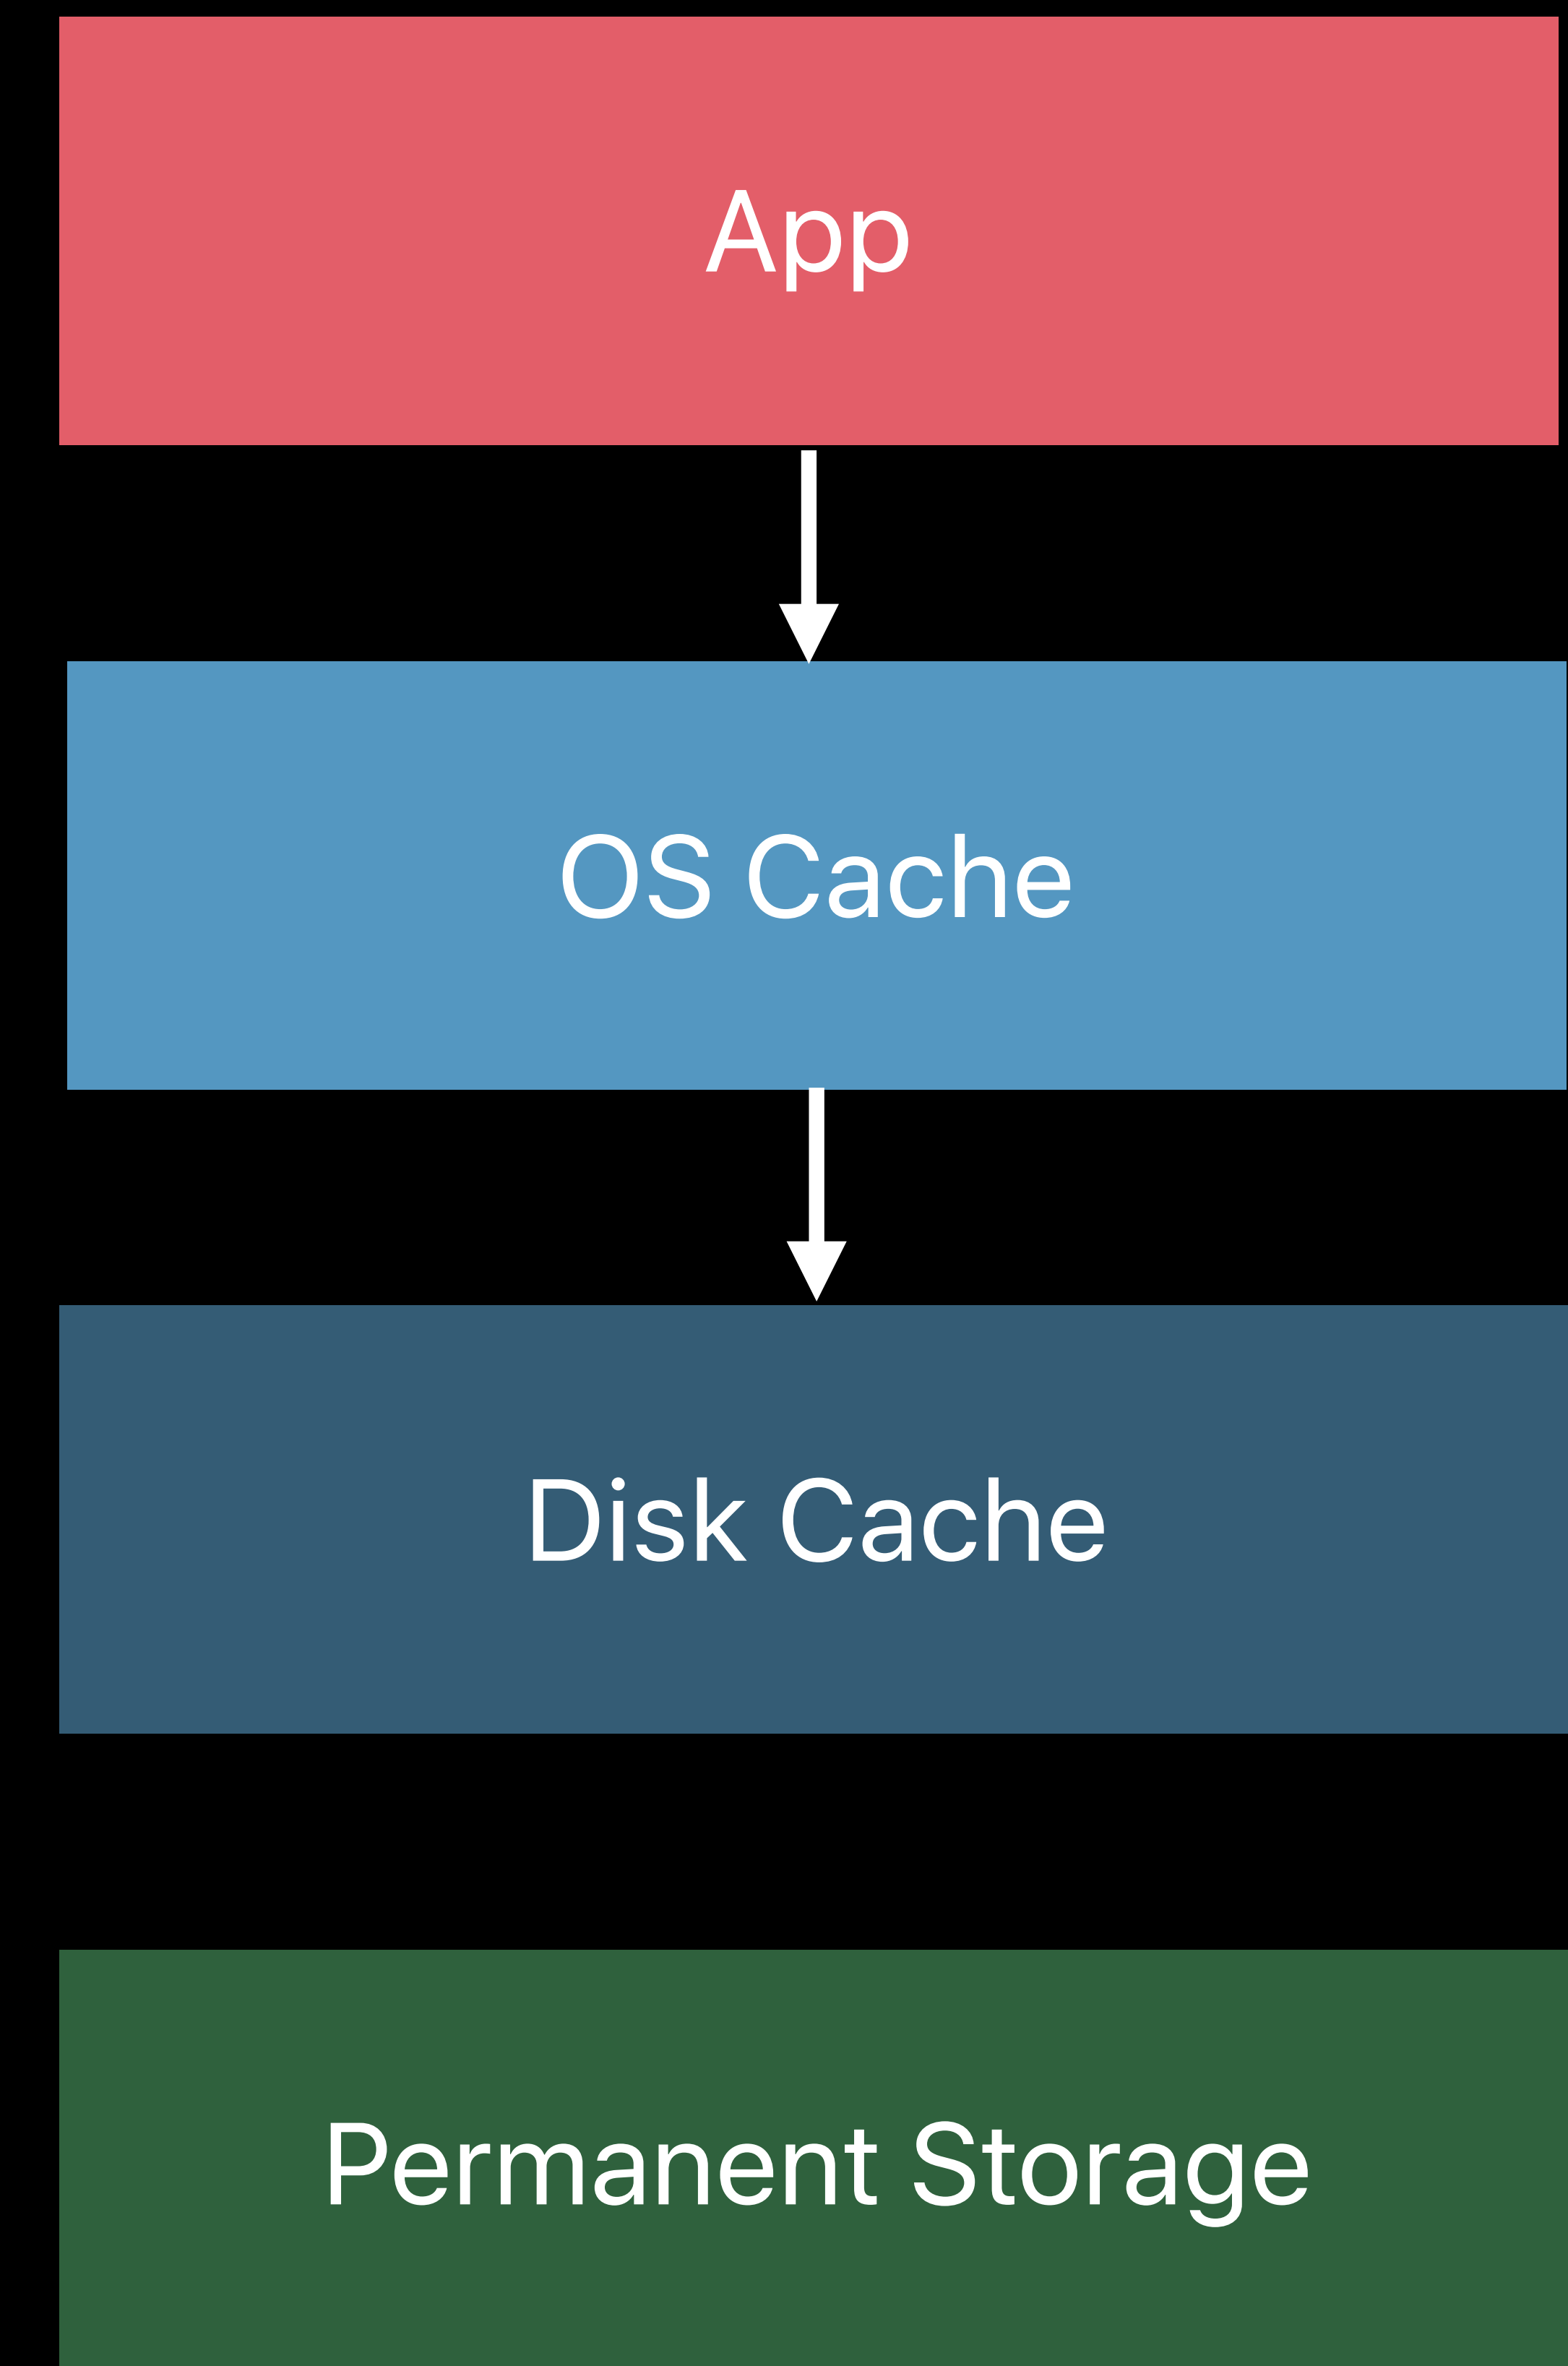




# fsync()

Flushes cached data to disk

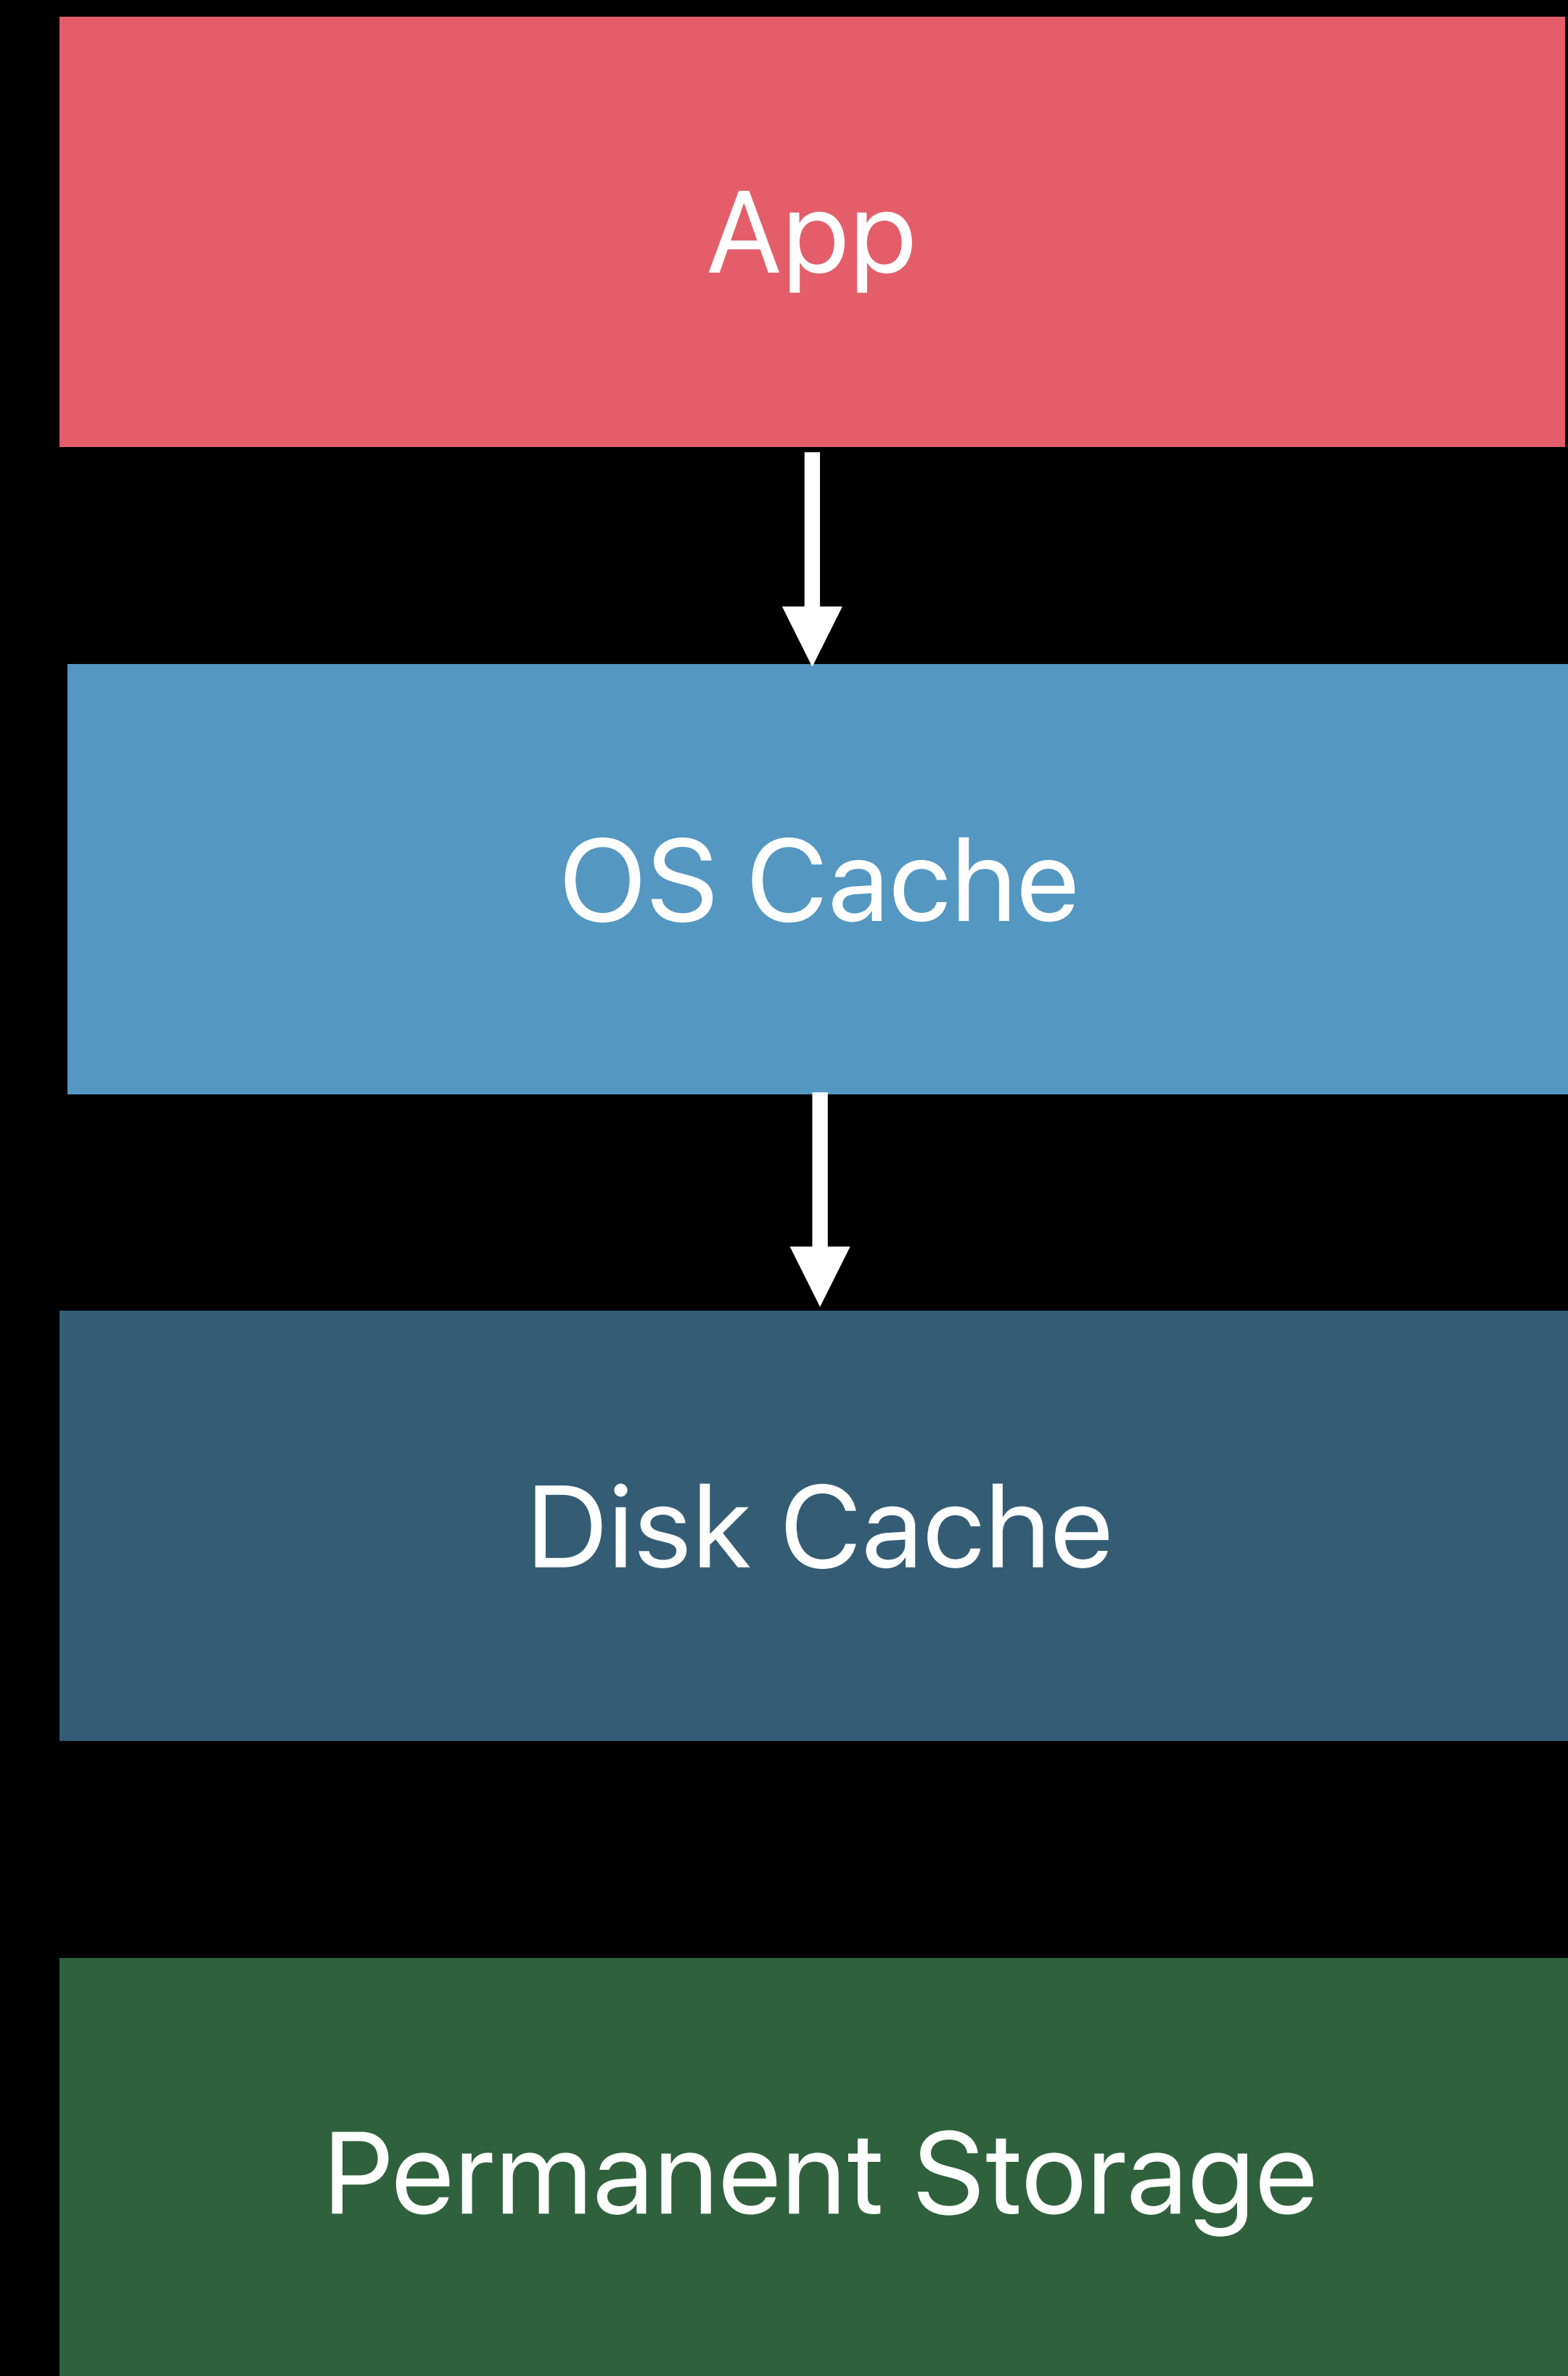
- Moves data from the OS cache to the disk cache
- Data may not be written to permanent storage immediately
- Does not guarantee write ordering
- Expensive



# fsync()

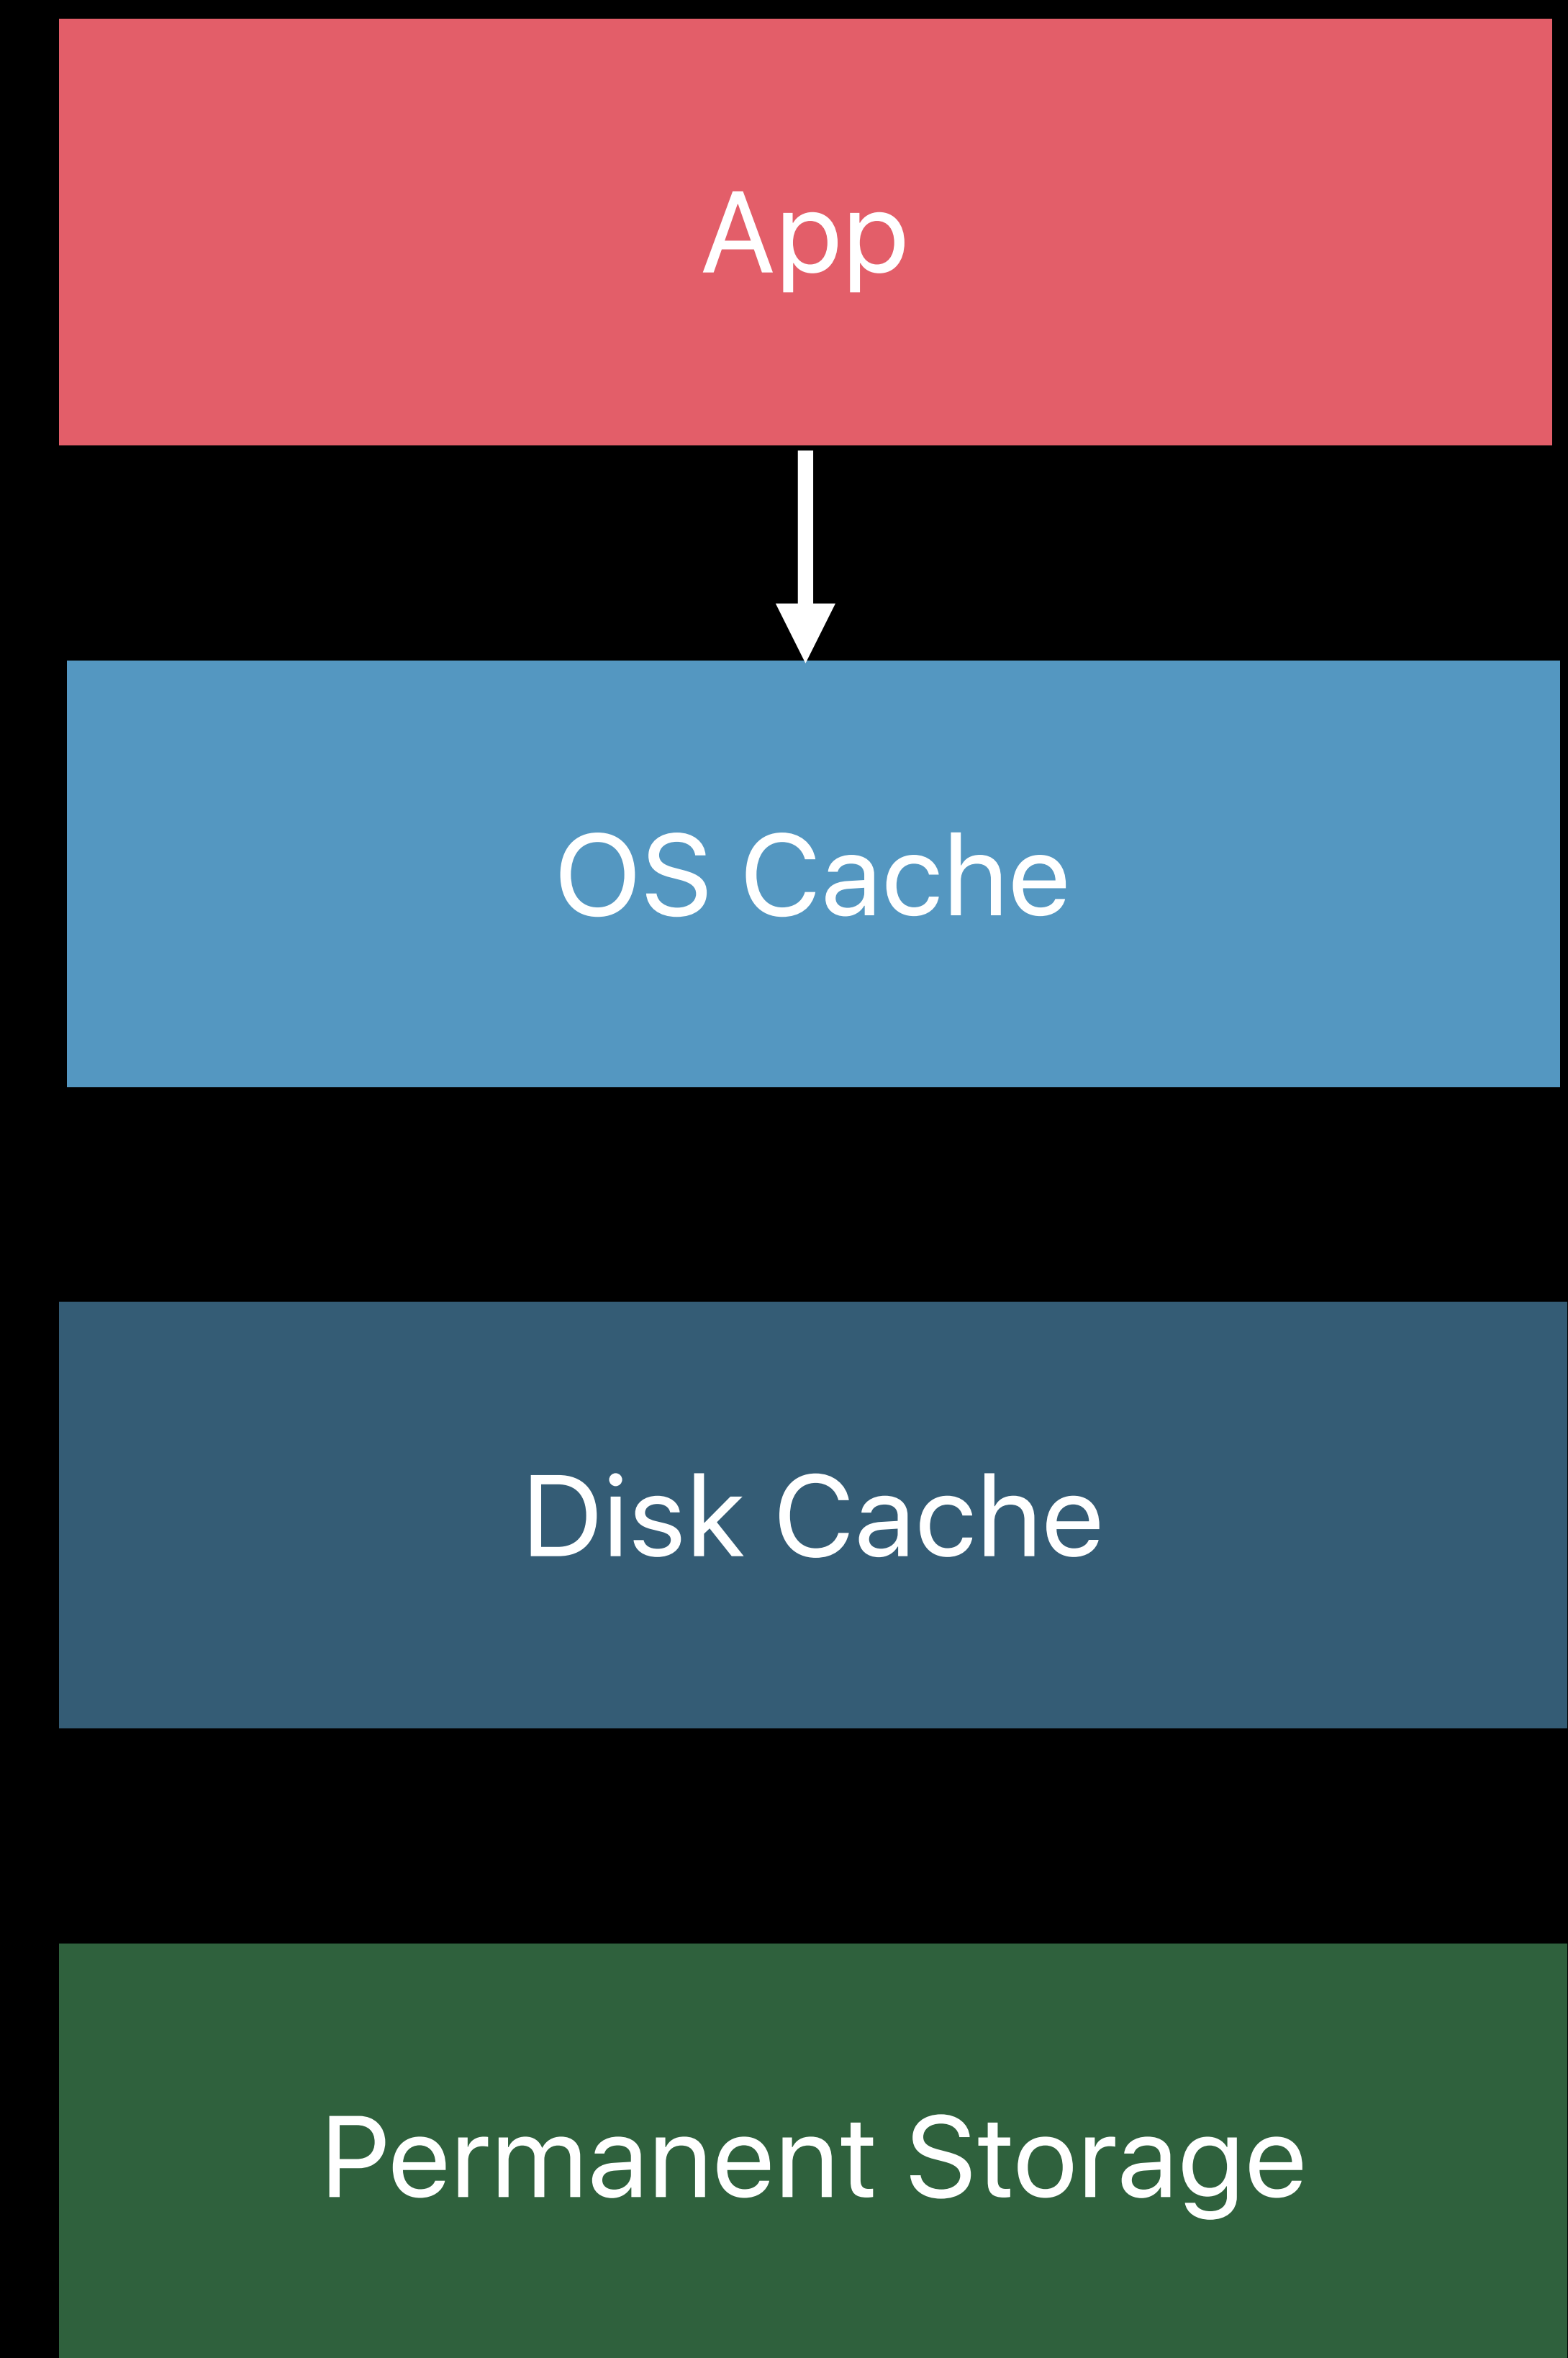
Flushes cached data to disk

- Moves data from the OS cache to the disk cache
- Data may not be written to permanent storage immediately
- Does not guarantee write ordering
- Expensive
- Done periodically by the OS



# F\_FULLFSYNC

Drive flush its cache to disk

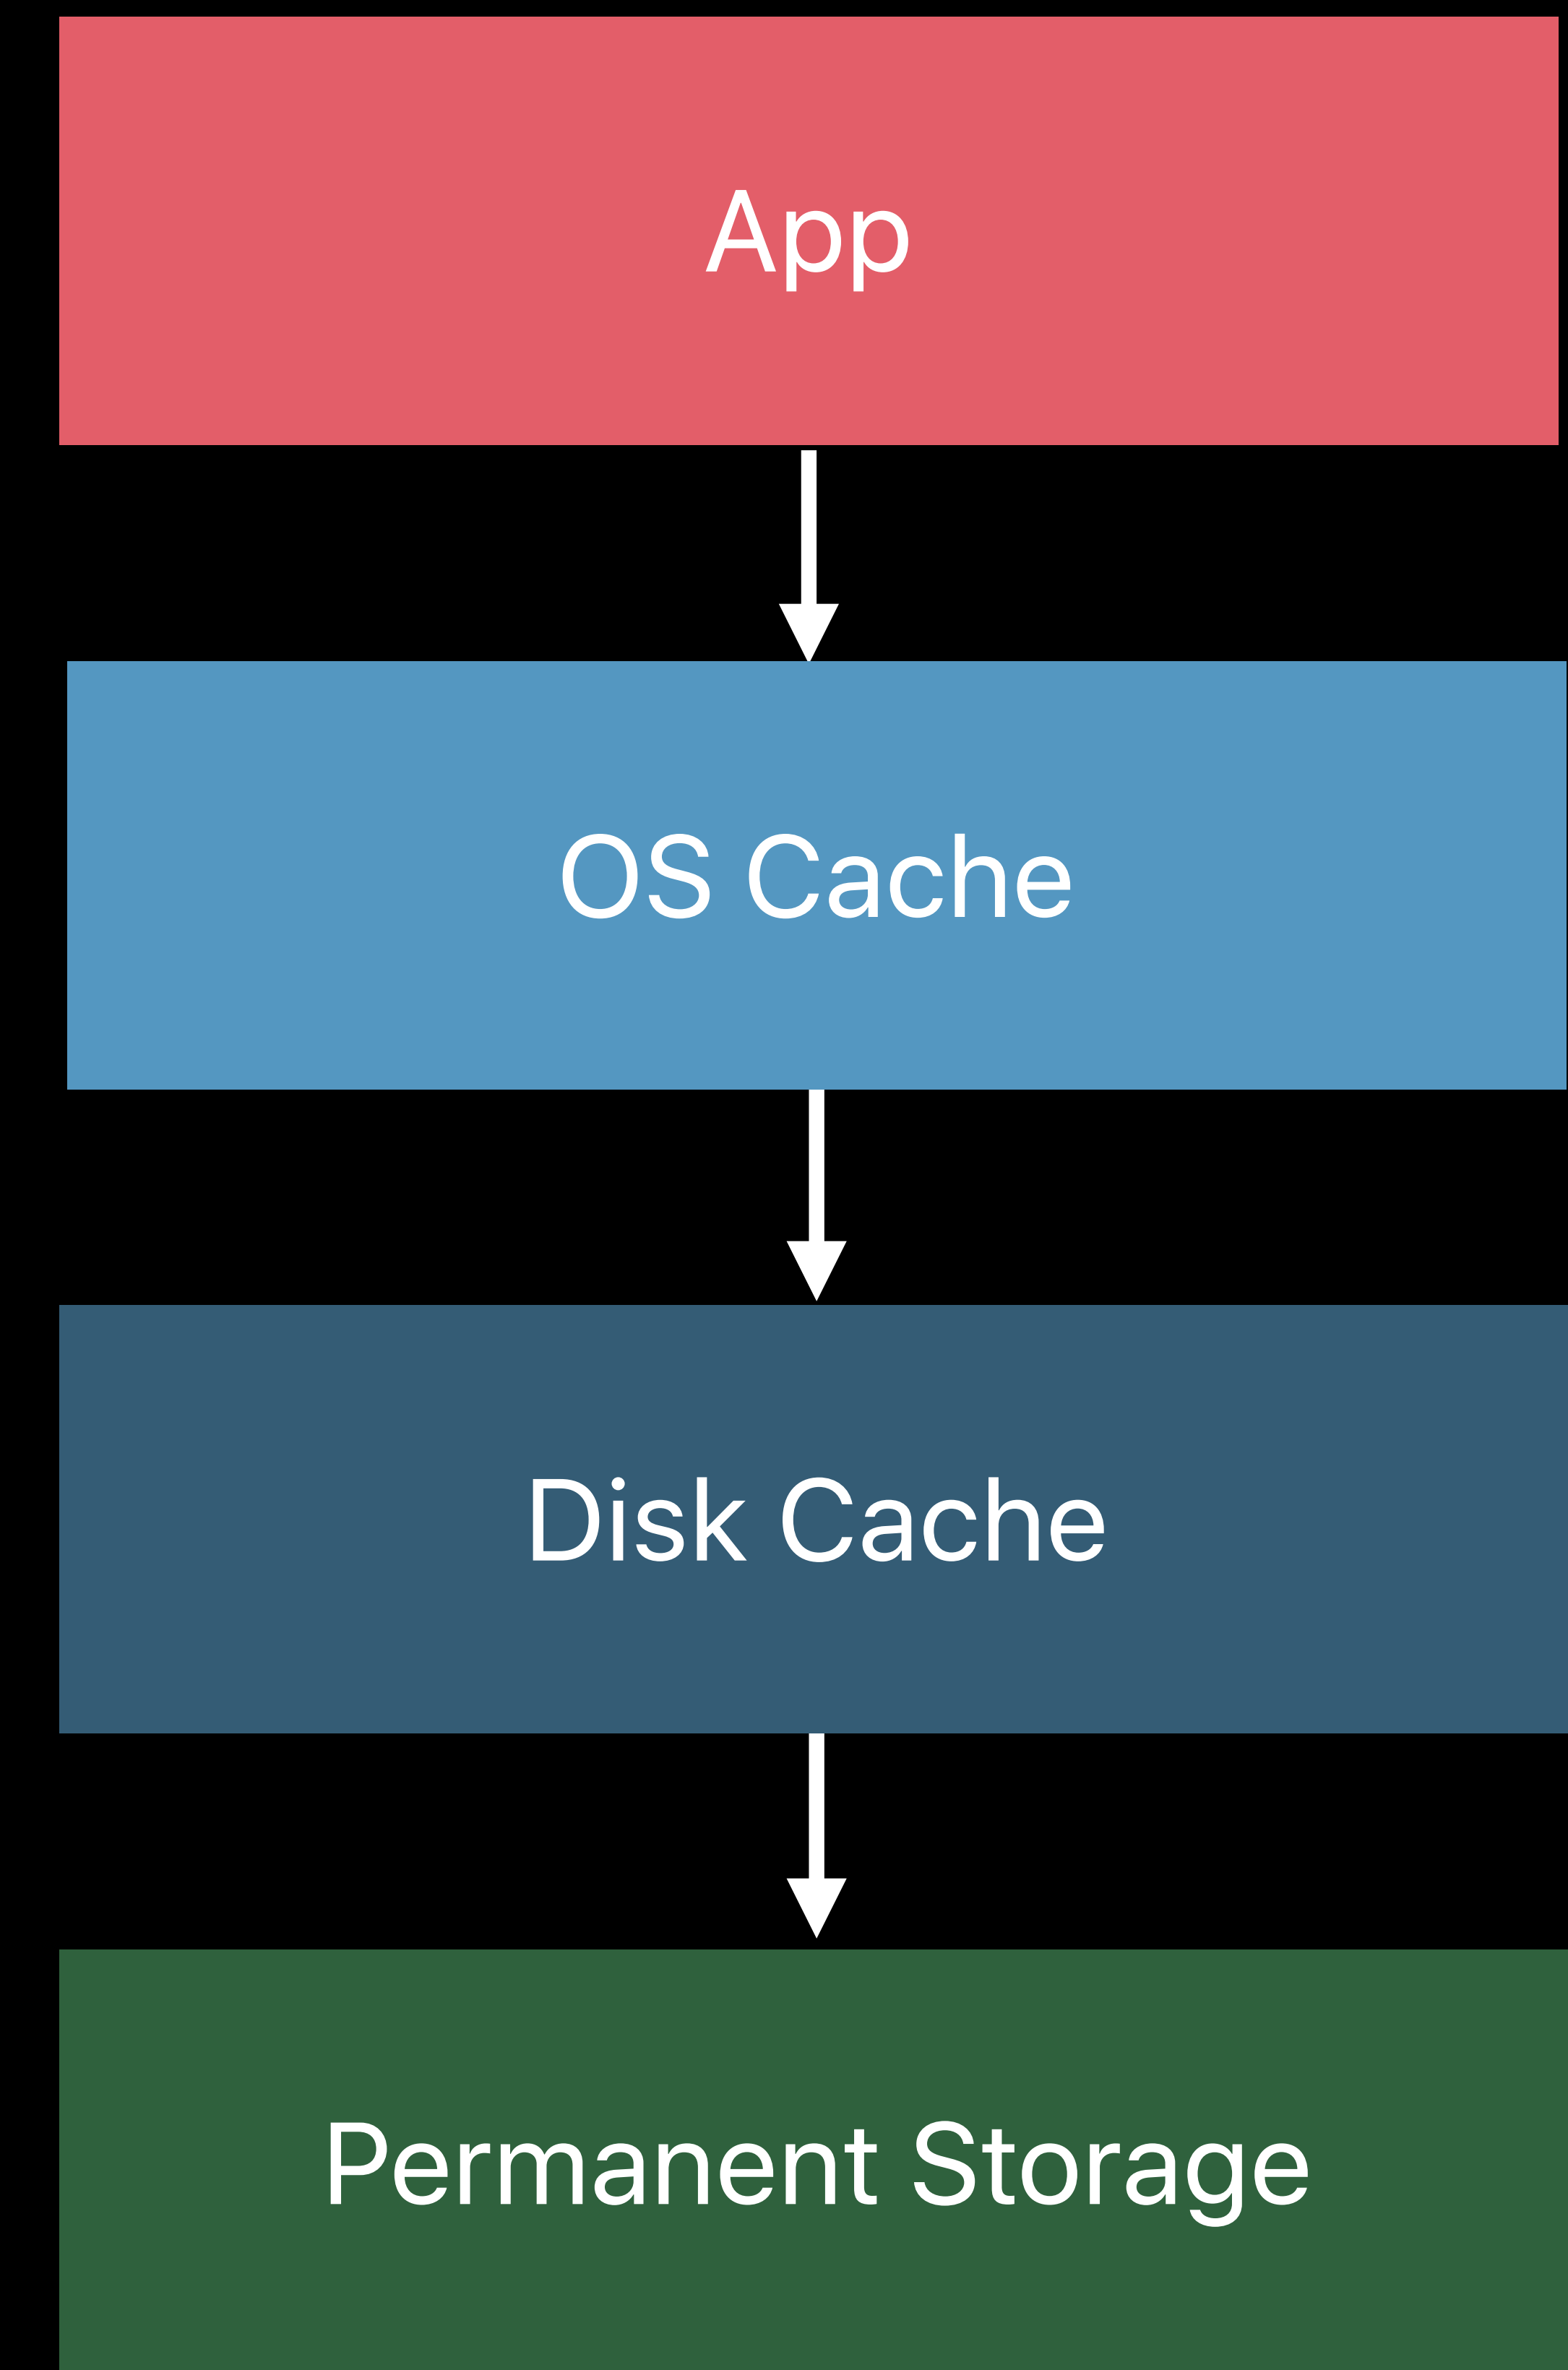




# F\_FULLFSYNC

Drive flush its cache to disk

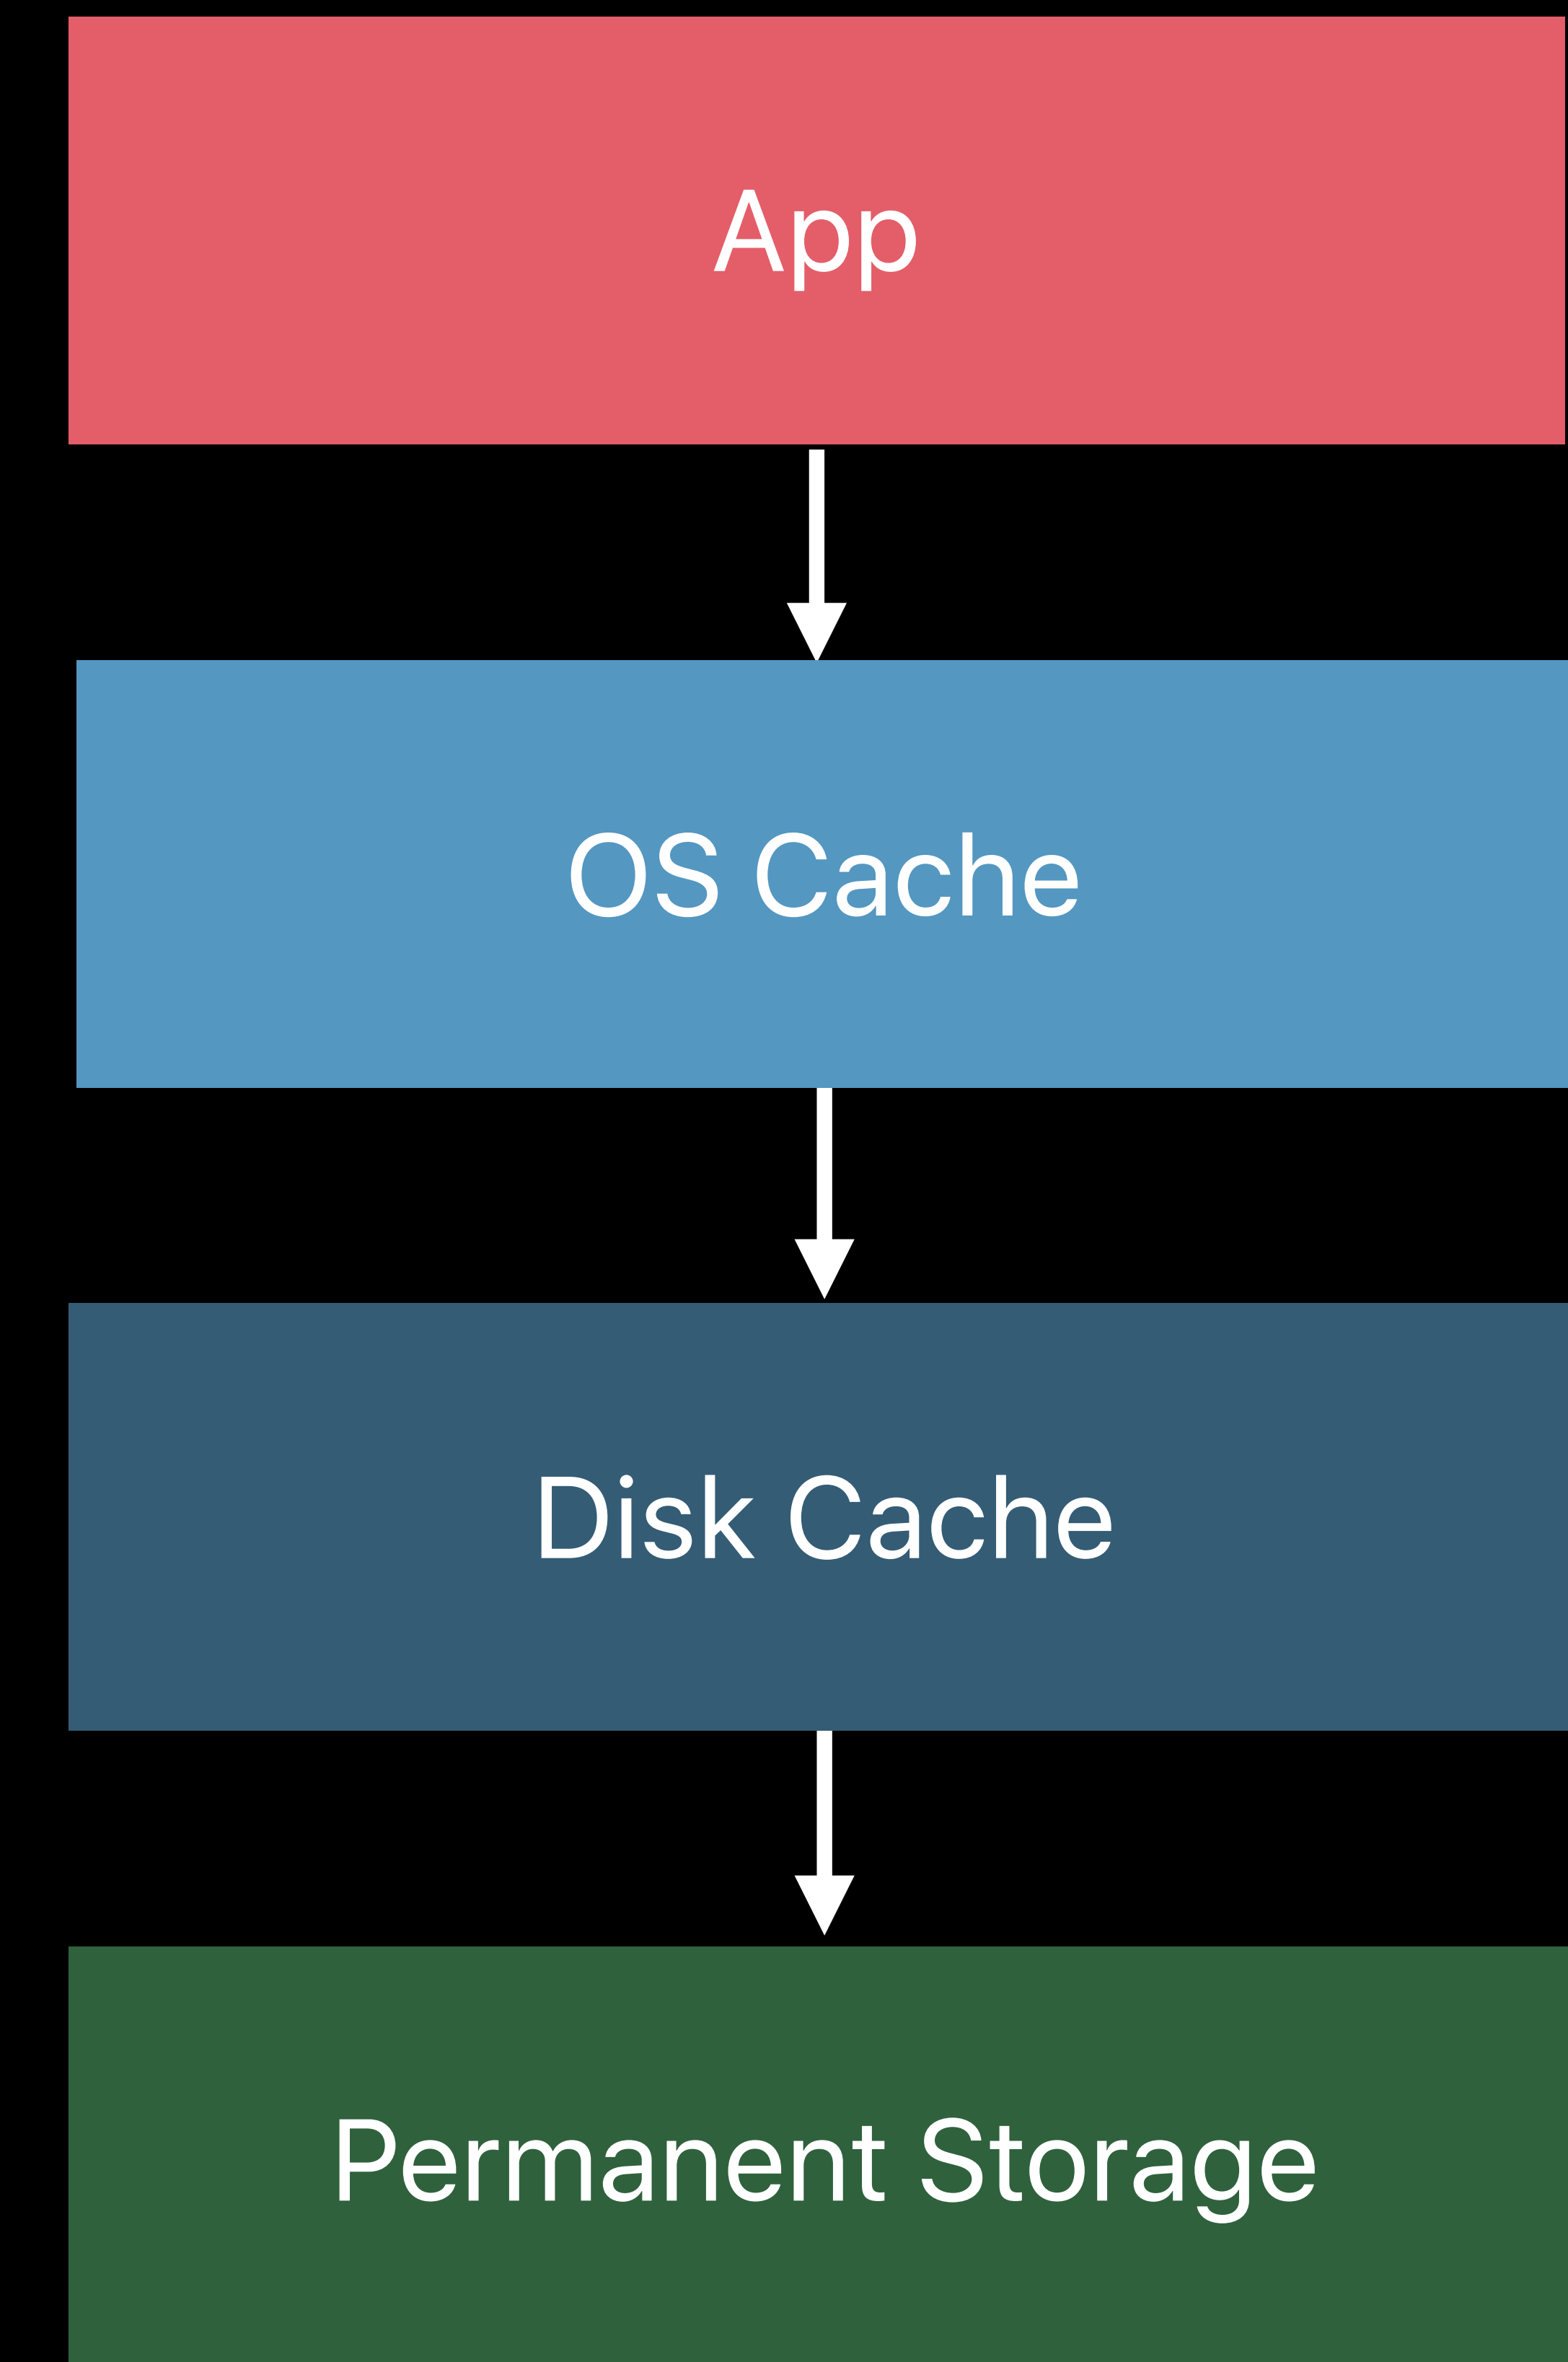
- Causes all data in disk cache to be flushed



# F\_FULLFSYNC

Drive flush its cache to disk

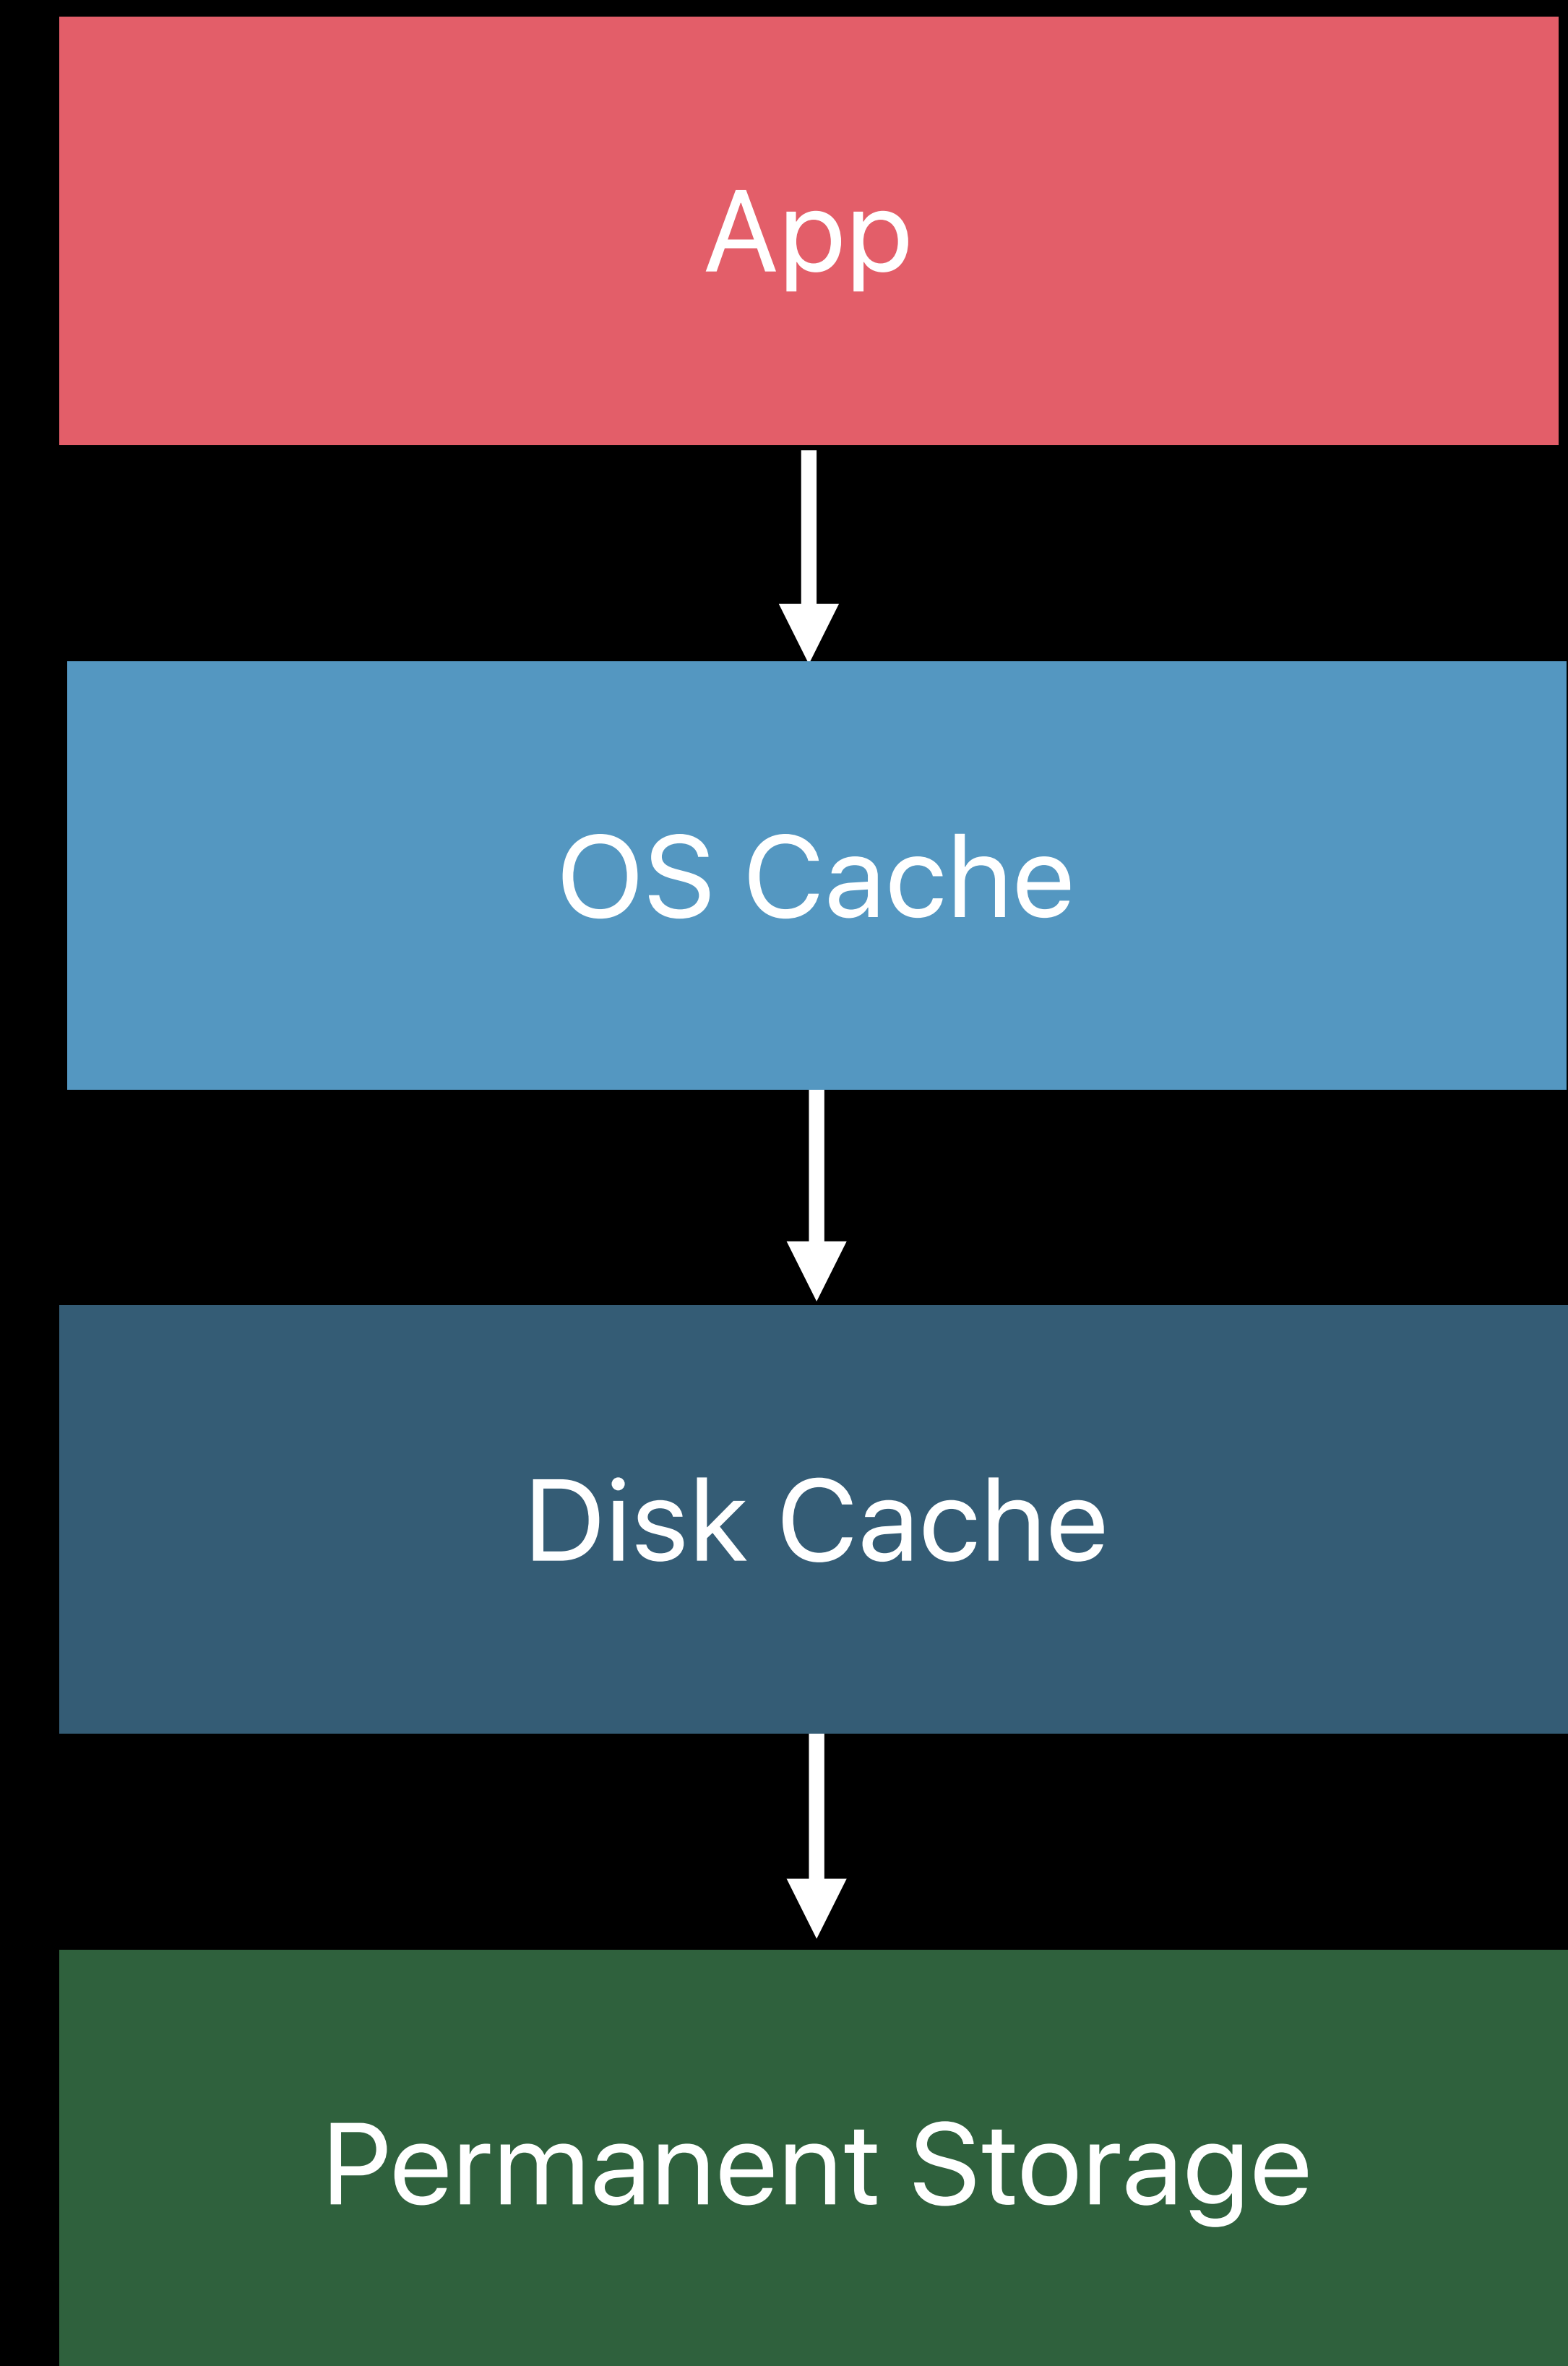
- Causes all data in disk cache to be flushed
- Expensive



# F\_FULLFSYNC

Drive flush its cache to disk

- Causes all data in disk cache to be flushed
- Expensive
- Done periodically by OS





# F\_BARRIERFSYNC

Enforces I/O ordering

# F\_BARRIERFSYNC

Enforces I/O ordering

- `fsync()` with a barrier

# F\_BARRIERFSYNC

Enforces I/O ordering

- `fsync()` with a barrier

Much less expensive than `F_FULLFSYNC`

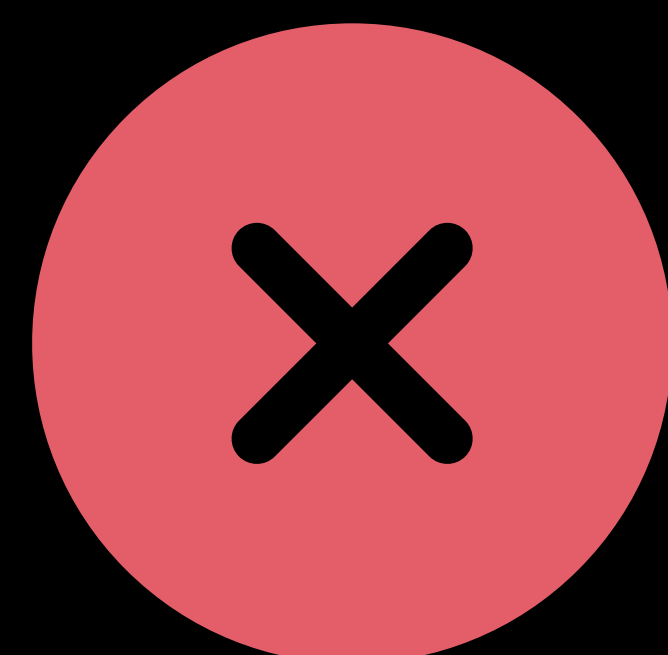


# Enforcing I/O Ordering



`F_BARRIERFSYNC`

Ensures I/O ordering for the specified data



`F_FULLFSYNC`

Indiscriminately commits all disk cache data to permanent storage

# Serialized Data Files

# Plists, XML, and JSON

The good

- Convenient

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST



# Plists, XML, and JSON

The good

- Convenient
- Excellent for infrequently written data

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST

# Plists, XML, and JSON

The good

- Convenient
- Excellent for infrequently written data
- Easy to parse

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST

# Plists, XML, and JSON

The trade off

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST



# Plists, XML, and JSON

The trade off

- Entire file must be rewritten for every change

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST

# Plists, XML, and JSON

The trade off

- Entire file must be rewritten for every change
- Scales poorly

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST

# Plists, XML, and JSON

## The trade off

- Entire file must be rewritten for every change
- Scales poorly
- Easy to misuse

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST



# Plists, XML, and JSON

## The trade off

- Entire file must be rewritten for every change
- Scales poorly
- Easy to misuse
- Metadata intensive

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST

# Plists, XML, and JSON

The trade off

- Entire file must be rewritten for every change
- Scales poorly
- Easy to misuse
- Metadata intensive

Not meant to be a database

Application Group	Mine
CFBundleDevelopmentRegion	English
CFBundleExecutable	MyApp
CFBundleIconFile	MyAppIcon.icns
CFBundleIdentifier	com.me.myApp
CFBundleName	My Application
CFBundleGetInfoString	Copyright 2014, Me
CFBundleShortVersionString	1.0
CFBundleVersion	1.0
NSMainNibFile	MainMenu
NSPrincipalClass	MyClass

PLIST



# NSDictionary

Create, read, and modify example

Filesystem Activity > Filesystem Events				
Start^	Duration	Process	Thread	Narrative
00:02.814.992	166.06 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed open on path plist_example.plist
00:02.815.177	153.14 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed write on fd 3 with a requested size of 484.46 KiB. Amount of bytes actually processed: 484.46 KiB
00:02.815.332	474.19 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed fsync on fd 3 and path (plist_example.plist)
00:02.815.810	15.83 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed close on fd 3 ( plist_example.plist )
00:02.816.509	25.57 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed open on path plist_example.plist
00:02.816.537	2.00 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed fstat64 on fd 3 and path (plist_example.plist)
00:02.816.561	296.53 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed read on fd 3 with a requested size of 484.46 KiB. Amount of bytes actually processed: 484.46 KiB
00:02.816.859	5.23 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed close on fd 3 ( plist_example.plist )
00:02.877.305	164.45 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed open on path plist_example.plist
00:02.877.473	135.80 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed write on fd 3 with a requested size of 484.46 KiB. Amount of bytes actually processed: 484.46 KiB
00:02.877.611	284.92 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed fsync on fd 3 and path (plist_example.plist)
00:02.877.897	9.04 μs	plist_example (31544)	plist_example (pid: 31544, tid: 0x606...	plist_example (31544) performed close on fd 3 ( plist_example.plist )



# NSDictionary

## Modify example

Disk Usage > Disk I/O Statistics > plist\_example.plist

Start	Total Lat...	Process	Thread	Operation	Path	Extra Flags	Bytes	Backtrace
00:02.815.378	409.48 μs	plist_example (31544)	Main Thread 0x60662ef	Data Write	plist_e...ple.plist	ASYNC	488.00 KiB	fsync
00:02.877.647	239.84 μs	plist_example (31544)	Main Thread 0x60662ef	Data Write	plist_e...ple.plist	ASYNC	488.00 KiB	_NSWriteDataToFileWithExtendedAttributes -[NSDictionary(NSDictionary) writeToFile:atomically:] writeDictToFile(dict:fileName:) dictionaryModify(file:) main start



# NSDictionary

## Modify example

Disk Usage > Disk I/O Statistics > plist\_example.plist

Start	Total Lat...	Process	Thread	Operation	Path	Extra Flags	Bytes	Backtrace
00:02.815.378	409.48 μs	plist_example (31544)	Main Thread 0x60662ef	Data Write	plist_e...ple.plist	ASYNC	488.00 KiB	fsync
00:02.877.647	239.84 μs	plist_example (31544)	Main Thread 0x60662ef	Data Write	plist_e...ple.plist	ASYNC	488.00 KiB	fsync _NSWriteDataToFileWithExtendedAttributes -[NSDictionary(NSDictionary) writeToFile:atomically:] writeDictToFile(dict:fileName:) dictionaryModify(file:) main start

**Core Data**



# Core Data

Core Data management

- Built on SQLite

# Core Data

## Core Data management

- Built on SQLite
- Manages object graphs and relationships

# Core Data

## Core Data management

- Built on SQLite
- Manages object graphs and relationships
- Change tracking and notifications



# Core Data

## Core Data management

- Built on SQLite
- Manages object graphs and relationships
- Change tracking and notifications
- Automatic version tracking and multi-writer conflict resolution
  - Automatic connection pooling

# Core Data

## Core Data management

- Built on SQLite
- Manages object graphs and relationships
- Change tracking and notifications
- Automatic version tracking and multi-writer conflict resolution
  - Automatic connection pooling
- CloudKit integration

# Core Data

## Core Data management

- Built on SQLite
- Manages object graphs and relationships
- Change tracking and notifications
- Automatic version tracking and multi-writer conflict resolution
  - Automatic connection pooling
- CloudKit integration (new with iOS 13)



# Core Data

## Core Data management

- Built on SQLite
- Manages object graphs and relationships
- Change tracking and notifications
- Automatic version tracking and multi-writer conflict resolution
  - Automatic connection pooling
- CloudKit integration (new with iOS 13)
- Live queries

# Core Data

## Core Data management

- Automatic memory management

# Core Data

## Core Data management

- Automatic memory management
- Statement aggregation in transactions



# Core Data

## Core Data management

- Automatic memory management
- Statement aggregation in transactions
- Schema migrations

# Core Data

## Core Data management

- Automatic memory management
- Statement aggregation in transactions
- Schema migrations
- Denormalization

# Core Data

## Core Data management

- Automatic memory management
- Statement aggregation in transactions
- Schema migrations
- Denormalization (new with iOS 13)



# Core Data

## Core Data management

- Automatic memory management
- Statement aggregation in transactions
- Schema migrations
- Denormalization (new with iOS 13)
- And so much more!

**50-70%**

Reduction in code to support the model layer

**SQLite**



Connections

Journaling

Transactions

File size and privacy

Partial indexes

# Connections

# Opening and Closing Connections

Can cause expensive operations!

- Consistency checking
- Journal recovery
- Journal checkpointing



# Opening and Closing Connections

Can cause expensive operations!

- Consistency checking
- Journal recovery
- Journal checkpointing

Recommended usage model

- Keep connections open as long as possible
- Close only when necessary
- Pool connections on multi-threaded processes

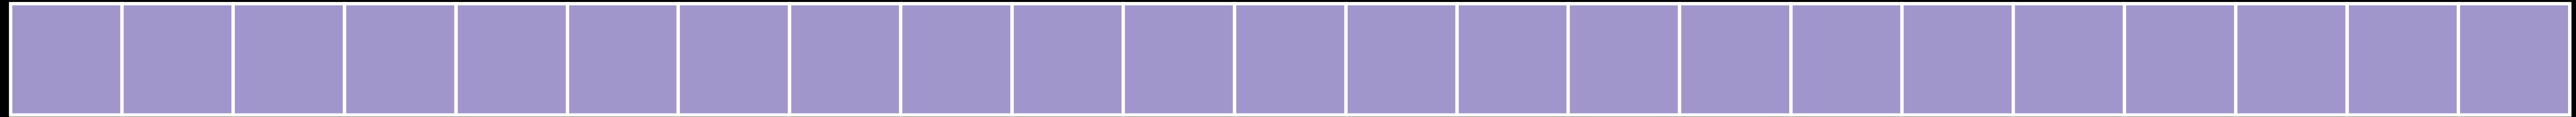
# Journaling

# Delete Mode Journaling

Default SQLite journaling mode

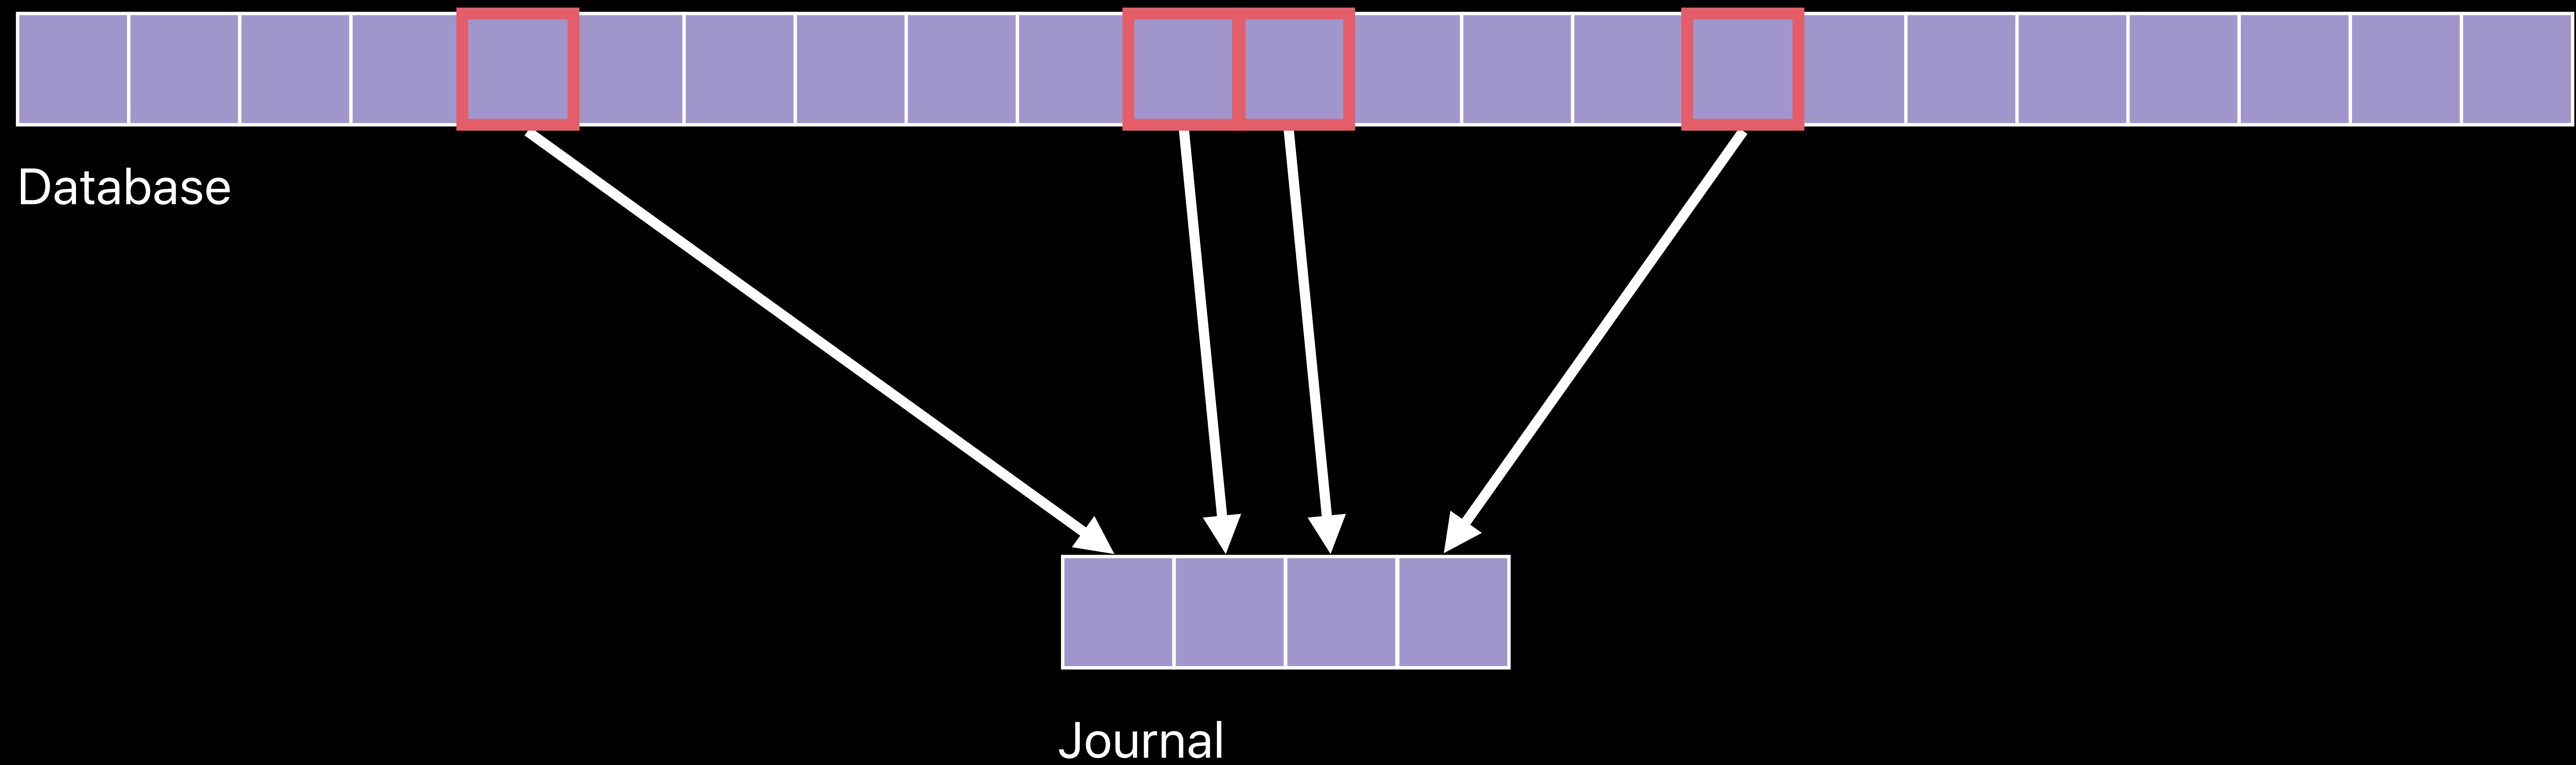


# How Does Delete Mode Journaling Work?

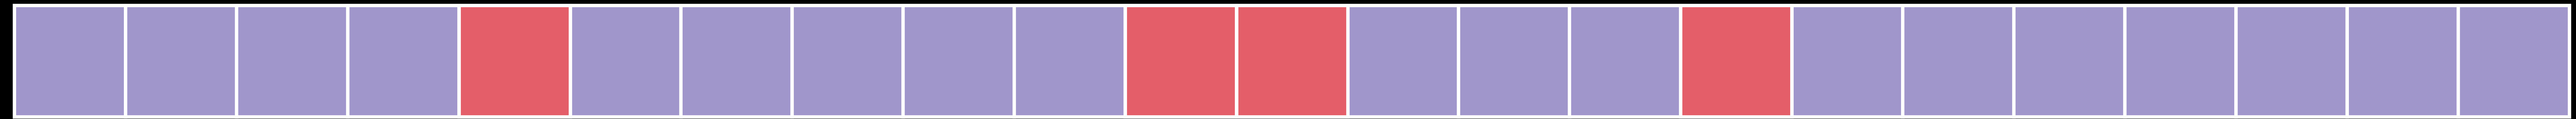


Database

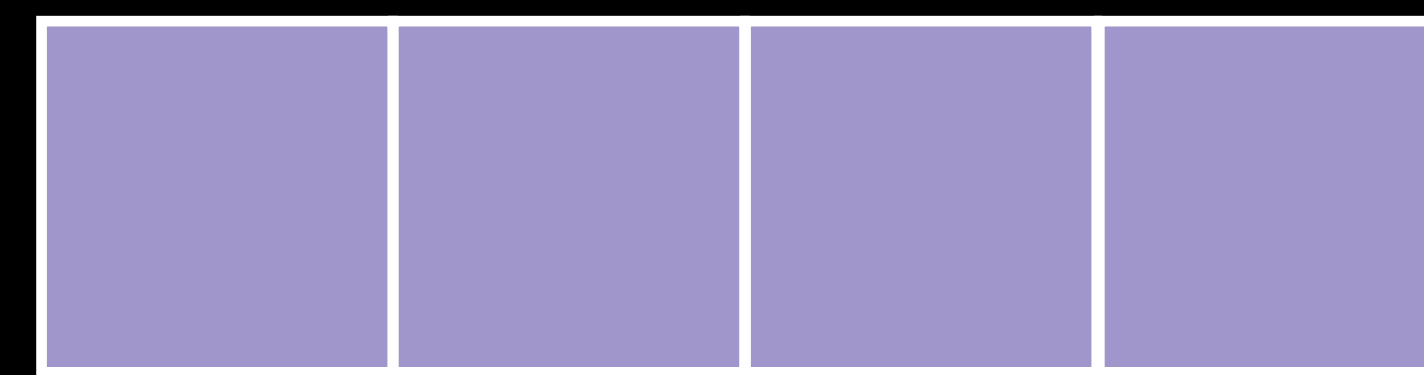
# How Does Delete Mode Journaling Work?



# How Does Delete Mode Journaling Work?



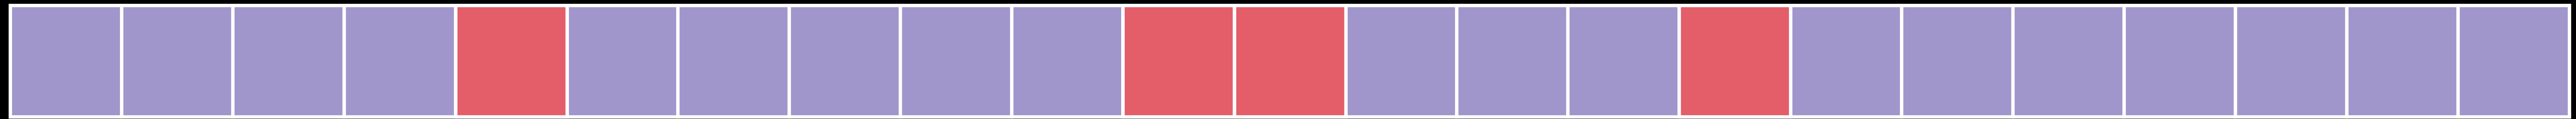
Database



Journal



# How Does Delete Mode Journaling Work?



Database

# WAL Mode Journaling Reduces Writes

```
PRAGMA journal_mode=WAL;
```

Write Ahead Logging

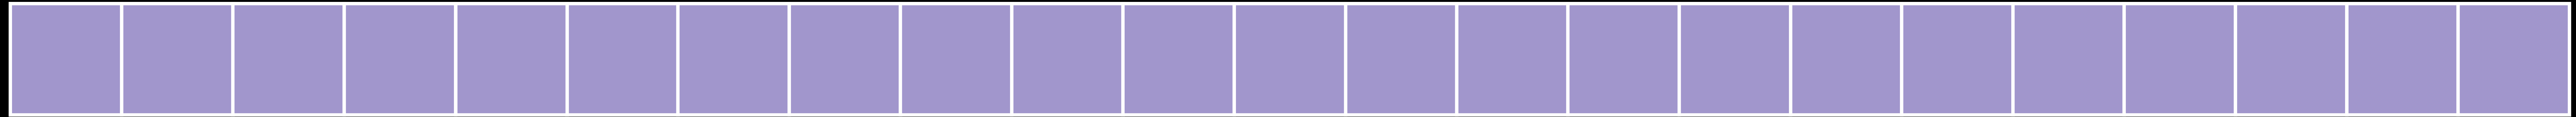
Combines multiple writes to the same page

Uses fewer barriers

Supports multiple readers at the same time as a writer

Supports snapshots

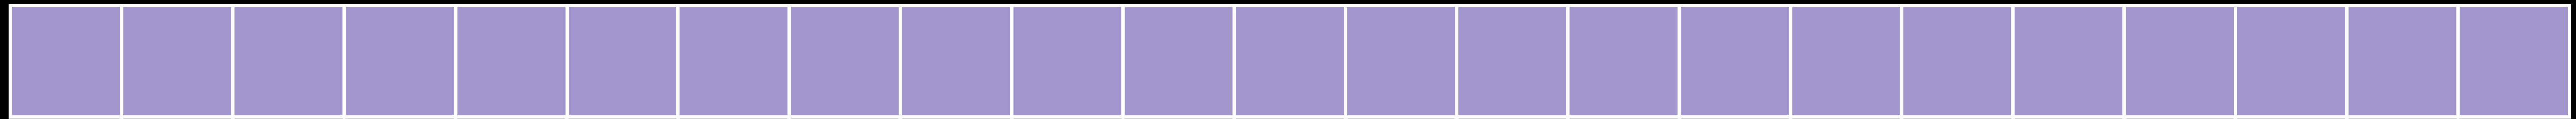
# How Does WAL Mode Journaling Work?



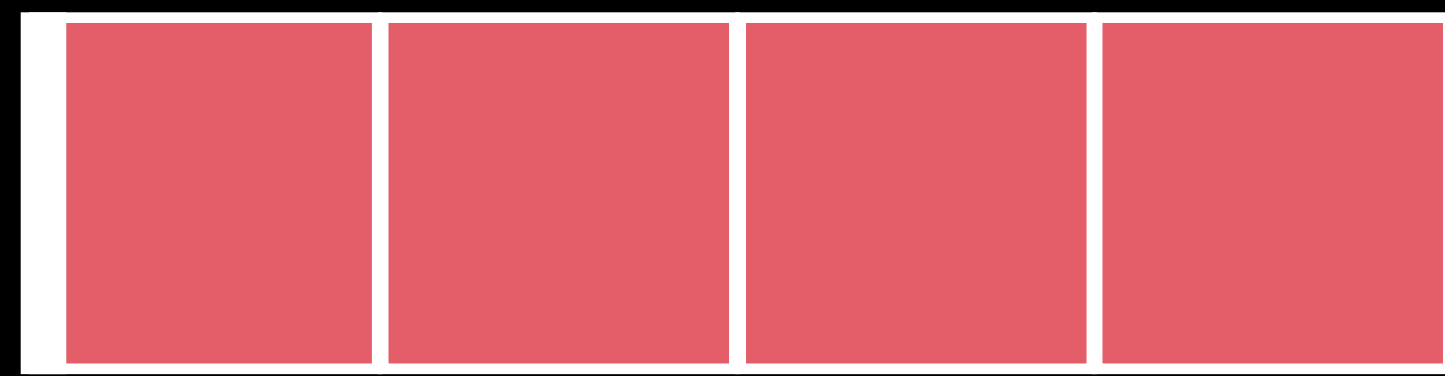
Database



# How Does WAL Mode Journaling Work?

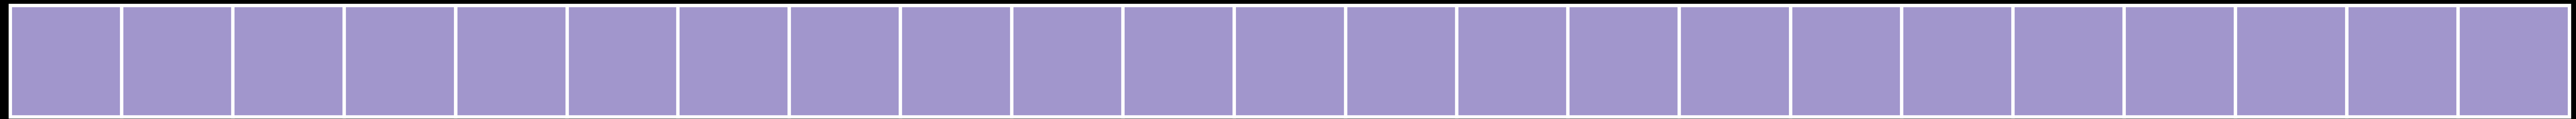


Database

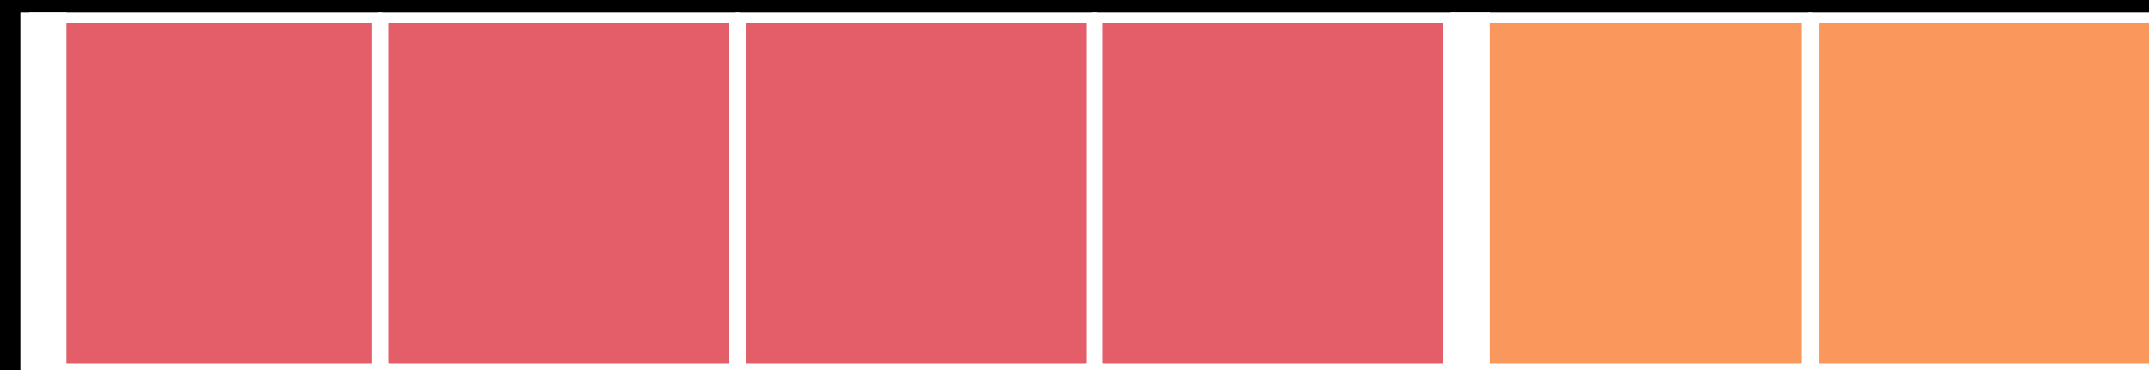


Write Ahead Log

# How Does WAL Mode Journaling Work?

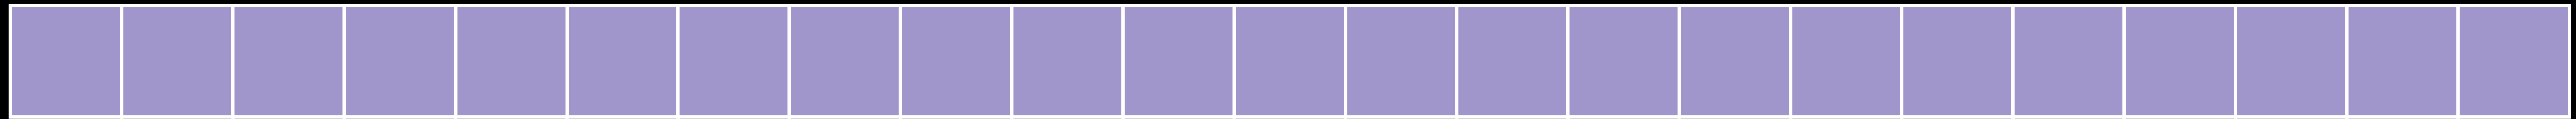


Database

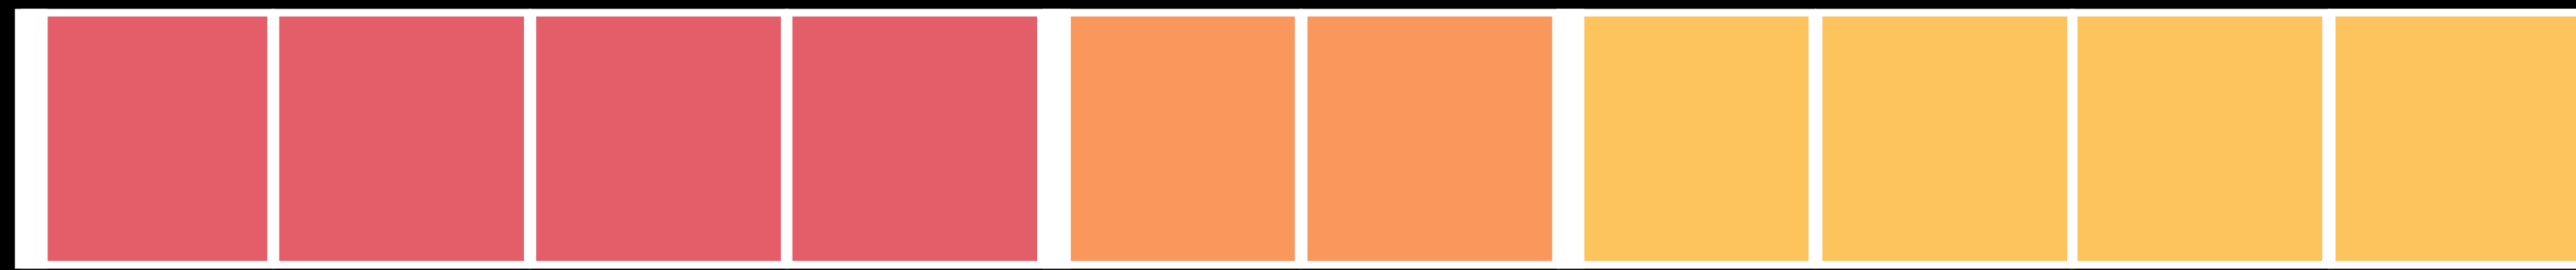


Write Ahead Log

# How Does WAL Mode Journaling Work?



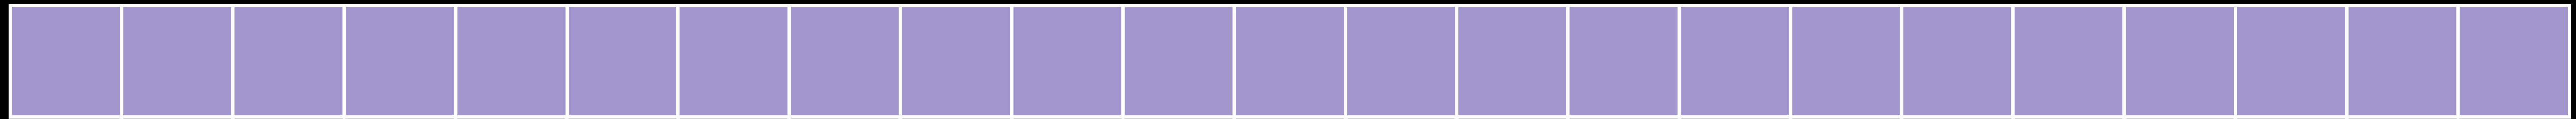
Database



Write Ahead Log



# How Does WAL Mode Journaling Work?

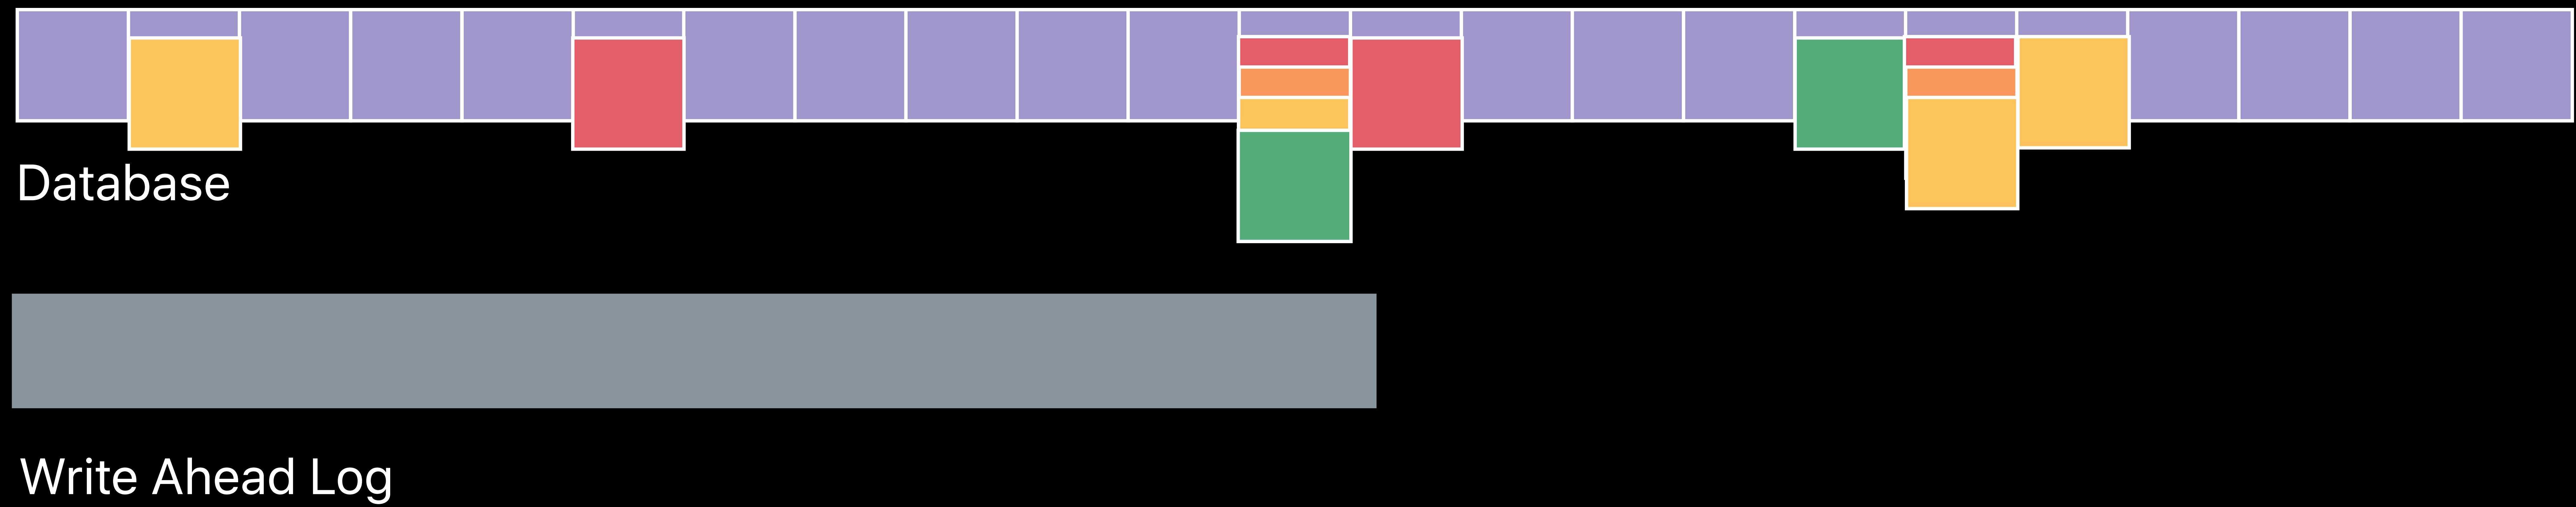


Database

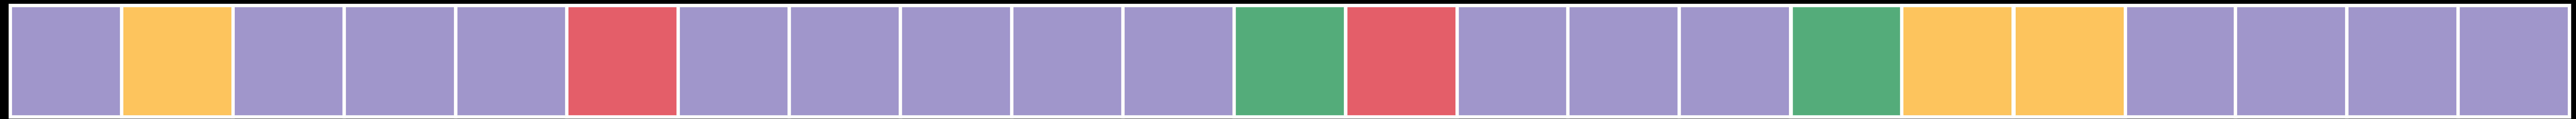


Write Ahead Log

# How Does WAL Mode Journaling Work?



# How Does WAL Mode Journaling Work?



Database



Write Ahead Log



WAL mode is more  
efficient for most use cases.

# Transactions

# Transactions Help Reduce Writes

Use for multiple `INSERT`, `UPDATE`, and `DELETE` statements

# Transactions Help Reduce Writes

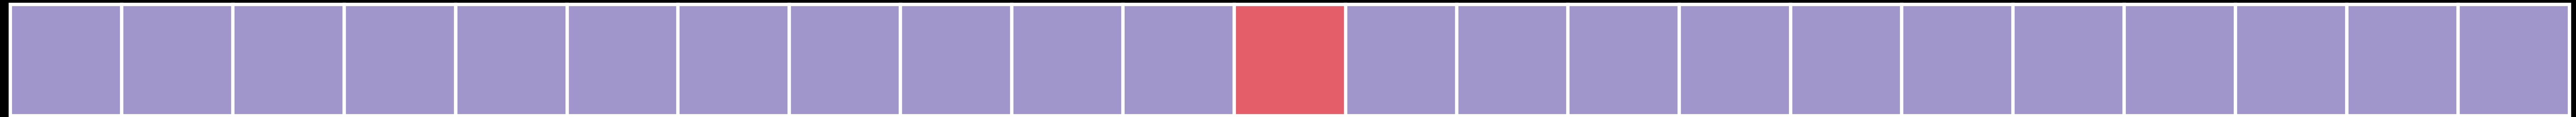
Use for multiple `INSERT`, `UPDATE`, and `DELETE` statements

Pages that are changed by multiple statements are only written once

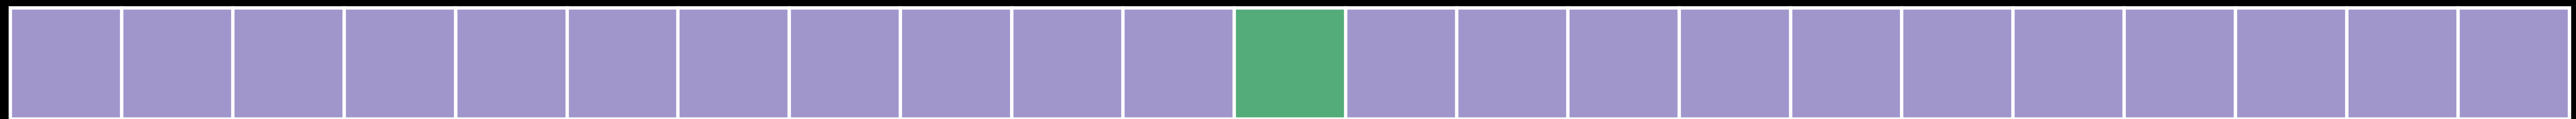


# Multiple Single Statement Transactions

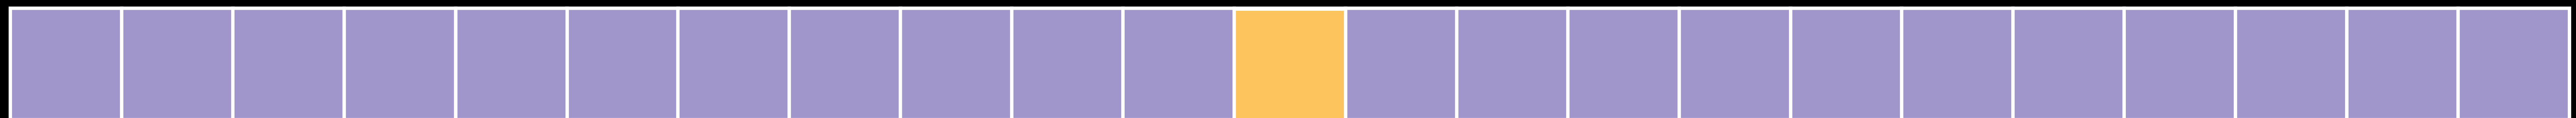
Transaction #1



Transaction #2

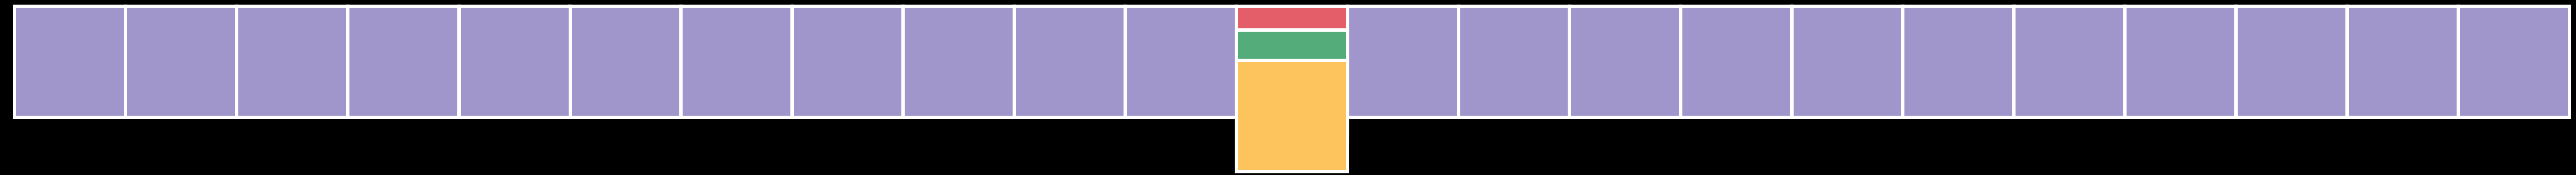


Transaction #3



# Multiple Statement Transaction

Transaction #1 with multiple statements



# Transactions Help Reduce Writes

Use for multiple `INSERT`, `UPDATE`, and `DELETE` statements

Pages that are changed by multiple statements are only written once

Useful for aggregating changes over time!

# File Size and Privacy



# File Size and Privacy

What happens when we delete data from a database?

Space containing the deleted data is marked as free

While no longer part of the database the deleted data is still on disk

# File Size and Privacy

How do we securely delete sensitive data?

# File Size and Privacy

```
PRAGMA schema.secure_delete=FAST;
```

How do we securely delete sensitive data?

# File Size and Privacy

```
PRAGMA schema.secure_delete=FAST;
```

How do we securely delete sensitive data?

- Automatically zeros deleted data
- No cost for data within the same page as the header
- Default behavior starting with iOS 13



# File Size and Privacy

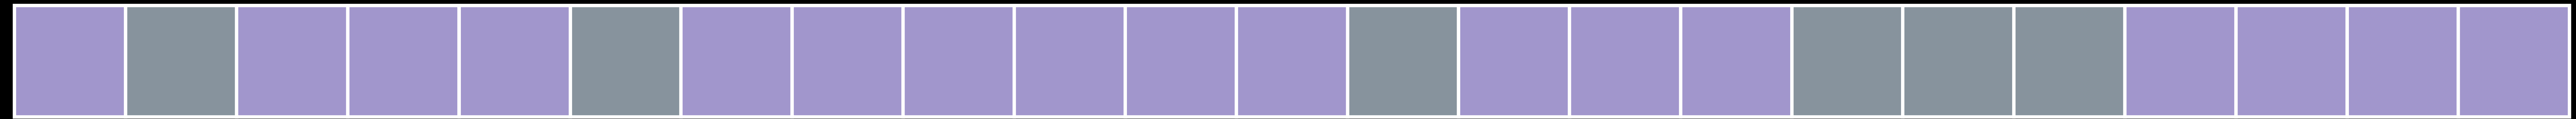
Don't use `VACUUM`

# File Size and Privacy

Don't use `VACUUM`

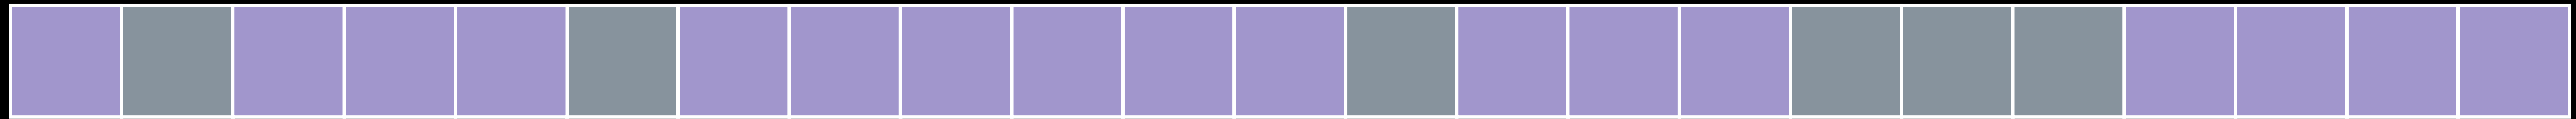
- `VACUUM` is a slow I/O intensive operation

# Why Is **VACUUM** Expensive?



Database

# Why Is **VACUUM** Expensive?



Database



Journal



# Why Is **VACUUM** Expensive?

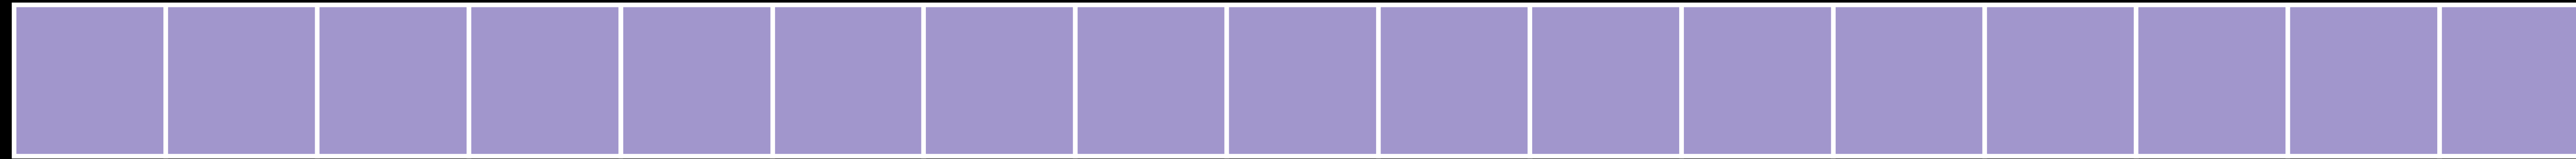


Database

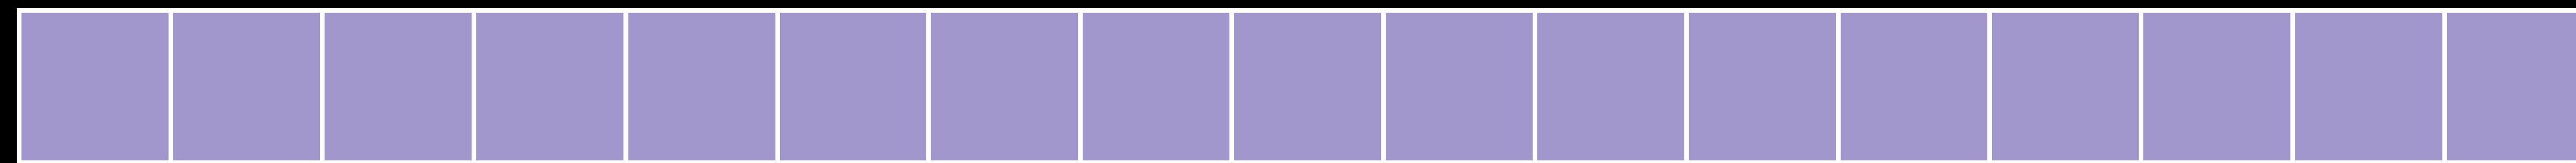


Journal

# Why Is **VACUUM** Expensive?

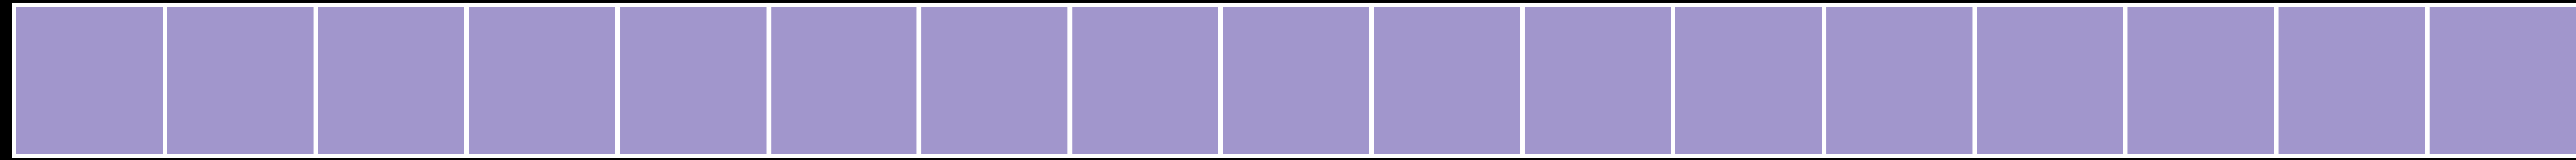


Database



Journal

# Why Is **VACUUM** Expensive?



Database

# File Size and Privacy

Don't use `VACUUM`

- `VACUUM` is a slow I/O intensive operation
- All valid data gets written at least twice!

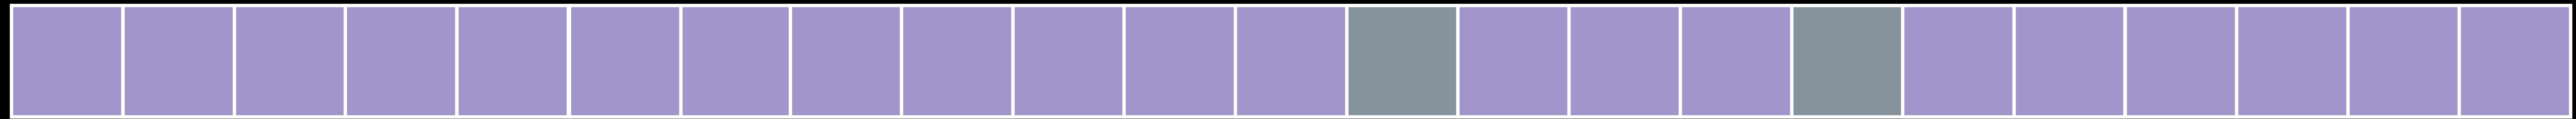


# File Size and Privacy

```
PRAGMA schema.auto_vacuum=INCREMENTAL;
```

```
PRAGMA schema.incremental_vacuum=(N);  
// 'N' is the number pages to be vacuumed
```

# How Is Incremental Auto Vacuum More Efficient?

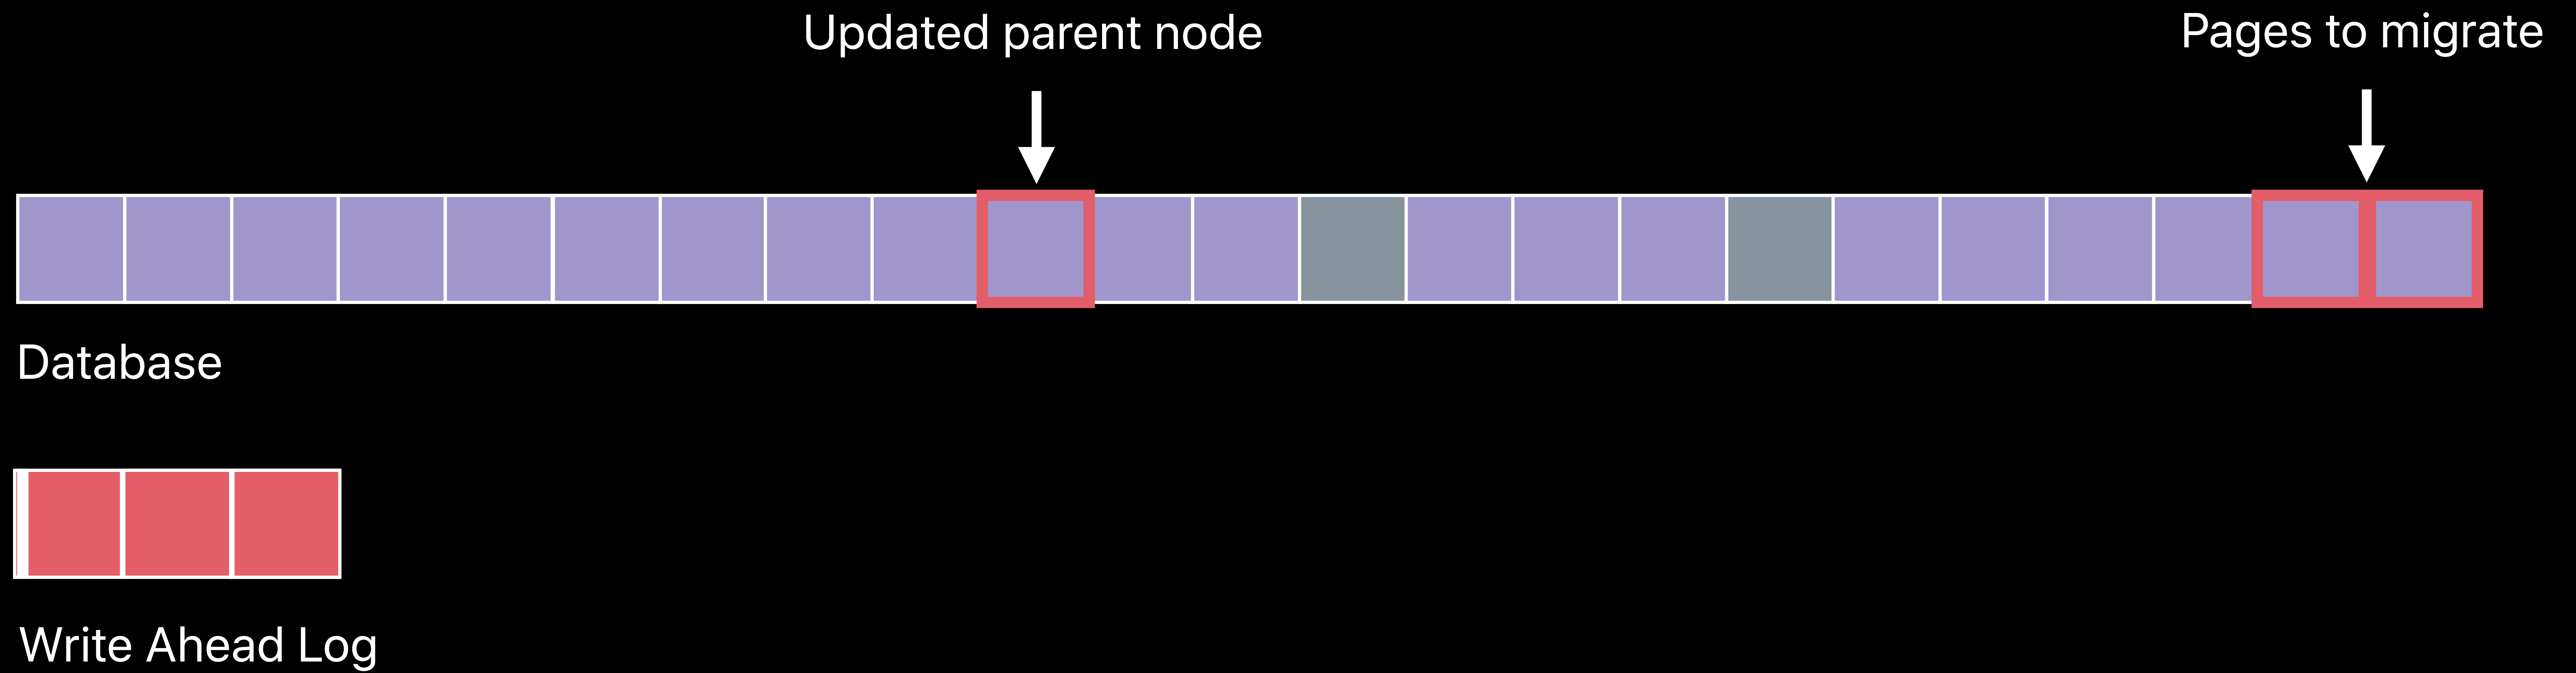


Database

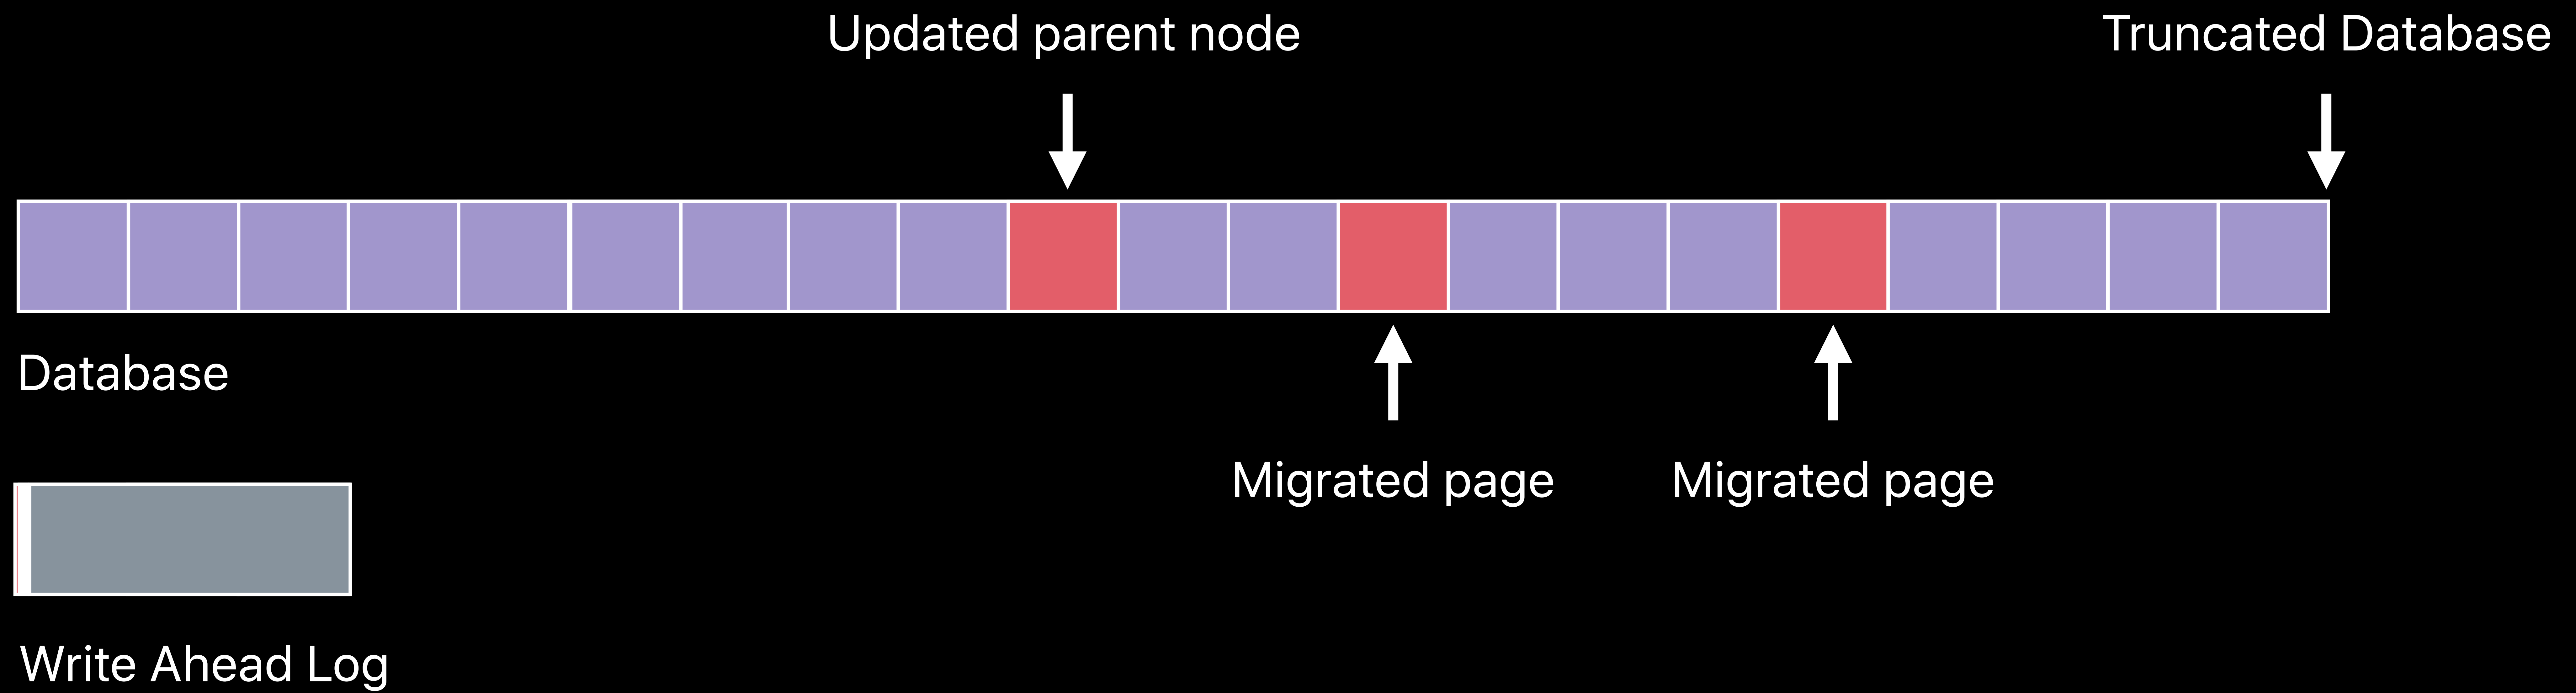


Write Ahead Log

# How Is Incremental Auto Vacuum More Efficient?



# How Is Incremental Auto Vacuum More Efficient?





Use incremental auto vacuum and  
fast secure delete to manage  
both file size and privacy.

# Partial Indexes

# Indexes

Faster `ORDER BY`, `GROUP BY`, and `WHERE` clauses

Each index adds additional I/O when writing to the database

# Partial Indexes

Indexes with a `WHERE` clause

Partial indexes only store data which matches the `WHERE` clause

Queries that do not match the `WHERE` clause cannot use the index



# SQLite Summary

# SQLite Summary

Best practices

Keep database connections open

# SQLite Summary

Best practices

Keep database connections open

Use WAL mode

# SQLite Summary

Best practices

Keep database connections open

Use WAL mode

Multiple statements per transaction



# SQLite Summary

Best practices

Keep database connections open

Use WAL mode

Multiple statements per transaction

Fast secure delete and auto vacuum incremental

# SQLite Summary

Best practices

Keep database connections open

Use WAL mode

Multiple statements per transaction

Fast secure delete and auto vacuum incremental

Partial indexes

# File Activity Instrument

Alejandro Lucena, Developer Tools

# File Activity Instrument



# File Activity Instrument

Supports all Apple devices

# File Activity Instrument

Supports all Apple devices

Support for tracing single processes and all processes

# File Activity Instrument

Supports all Apple devices

Support for tracing single processes and all processes

Obtains both logical and physical I/O information

# File Activity Instrument

Supports all Apple devices

Support for tracing single processes and all processes

Obtains both logical and physical I/O information

Offers automated reasoning



# Automated Reasoning

# Automated Reasoning

Excessive physical writes

# Automated Reasoning

Excessive physical writes

Failed I/O related calls

# Automated Reasoning

Excessive physical writes



















Failed I/O related calls


Suboptimal caching



# File Activity Instrument

Standard Custom Recent Filter

 Blank	 Activity Monitor	 Allocations	 App Launch	 Core Animation	 Core Data
 Counters	 Energy Log	 File Activity	 Game Performance	 Leaks	 Metal System Trace
					

 **File Activity**  
This template records and analyzes filesystem and disk I/O activity. Recommendations are generated if the template detects behavior that is symptomatic of general I/O anti-patterns.



# File Activity Instrument

The screenshot displays the Instruments application window. At the top, the title bar reads "Instruments" and the status bar shows "No Runs". The main interface includes a "Track Filter" section with a "Track Filter" button and an "All Tracks" button. Below this is a timeline with markers at 00:00.000, 00:10.000, 00:20.000, 00:30.000, and 00:40.000. The left sidebar lists four instruments: "Filesystem Suggestions" (highlighted in blue), "Filesystem Activity", "Disk Usage", and "Disk I/O Latency". Each instrument has a red circular icon and a label "Instrument" below it. The right pane shows a dark area with a faint blue line graph, likely representing the selected instrument's data.

Instruments

Alejandro's iMac Pro > All Processes

No Runs

Track Filter

All Tracks

00:00.000 00:10.000 00:20.000 00:30.000 00:40.000

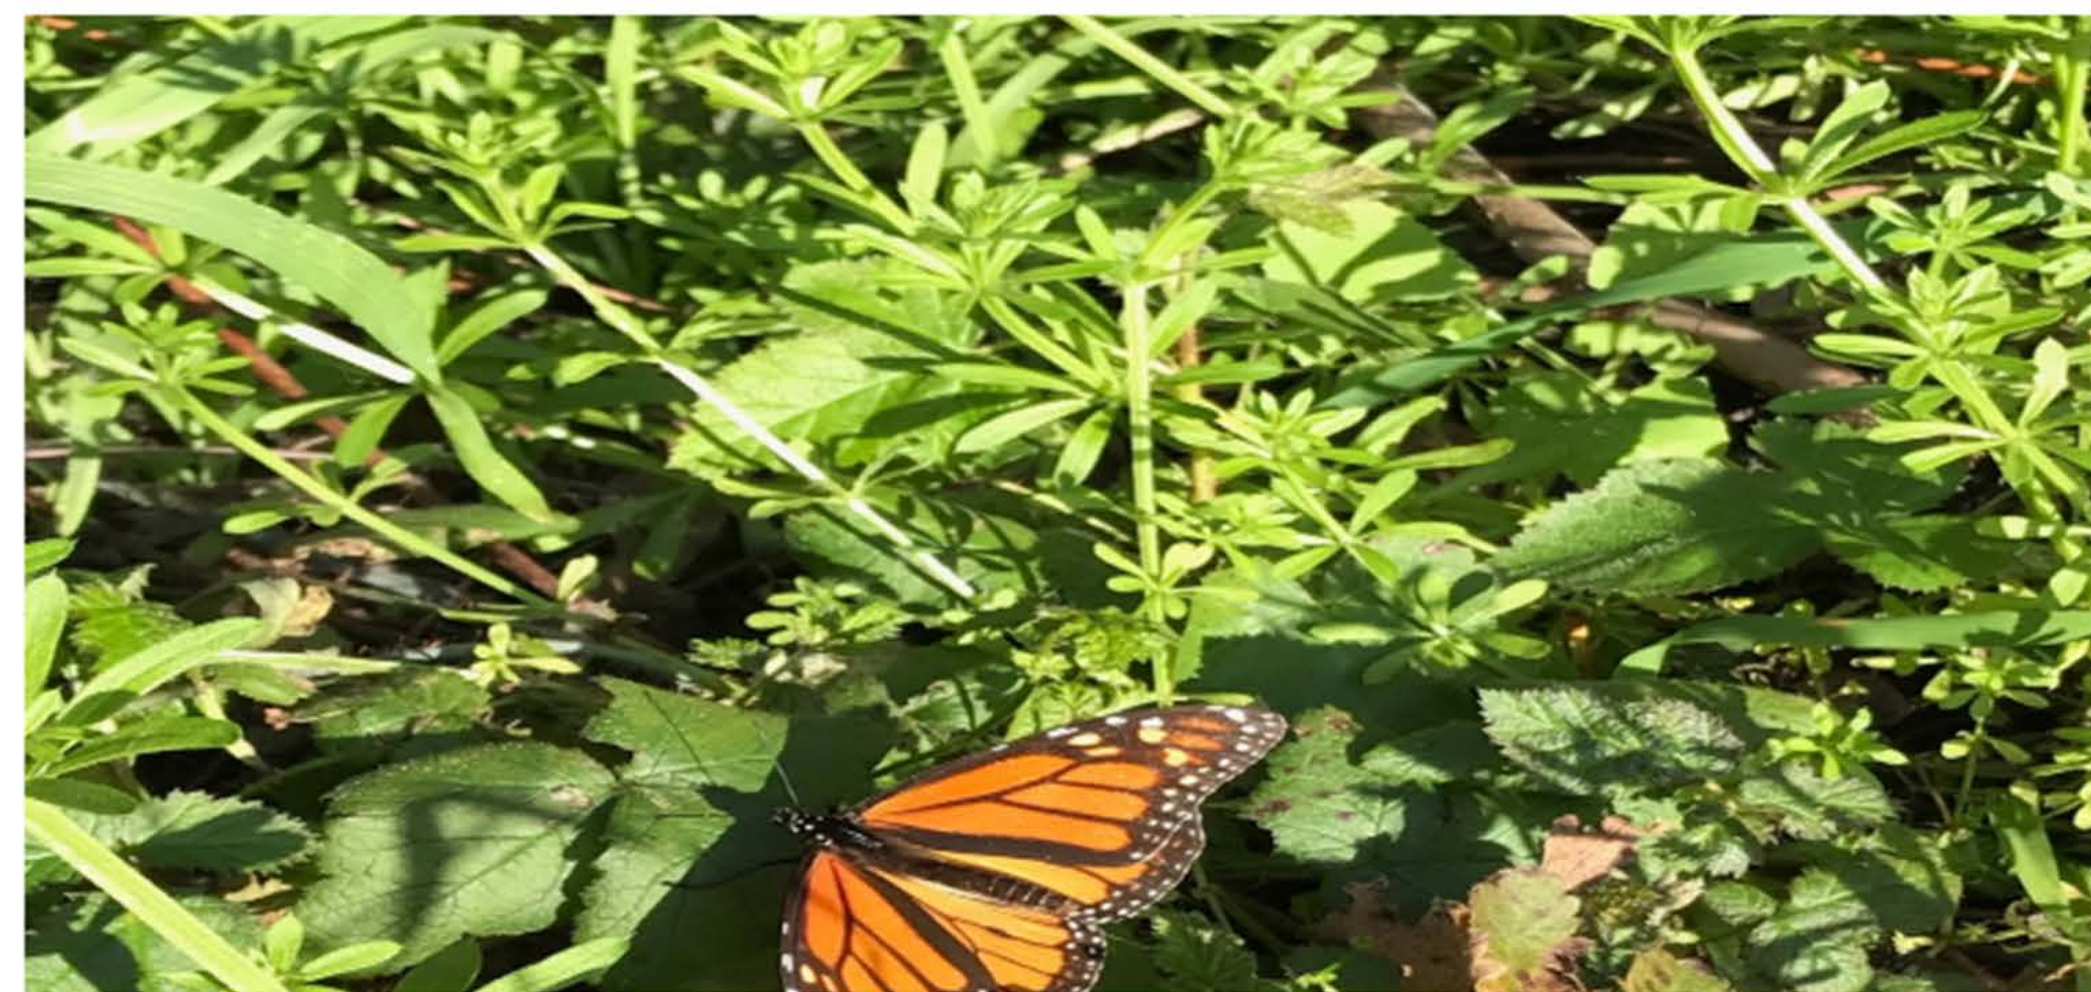
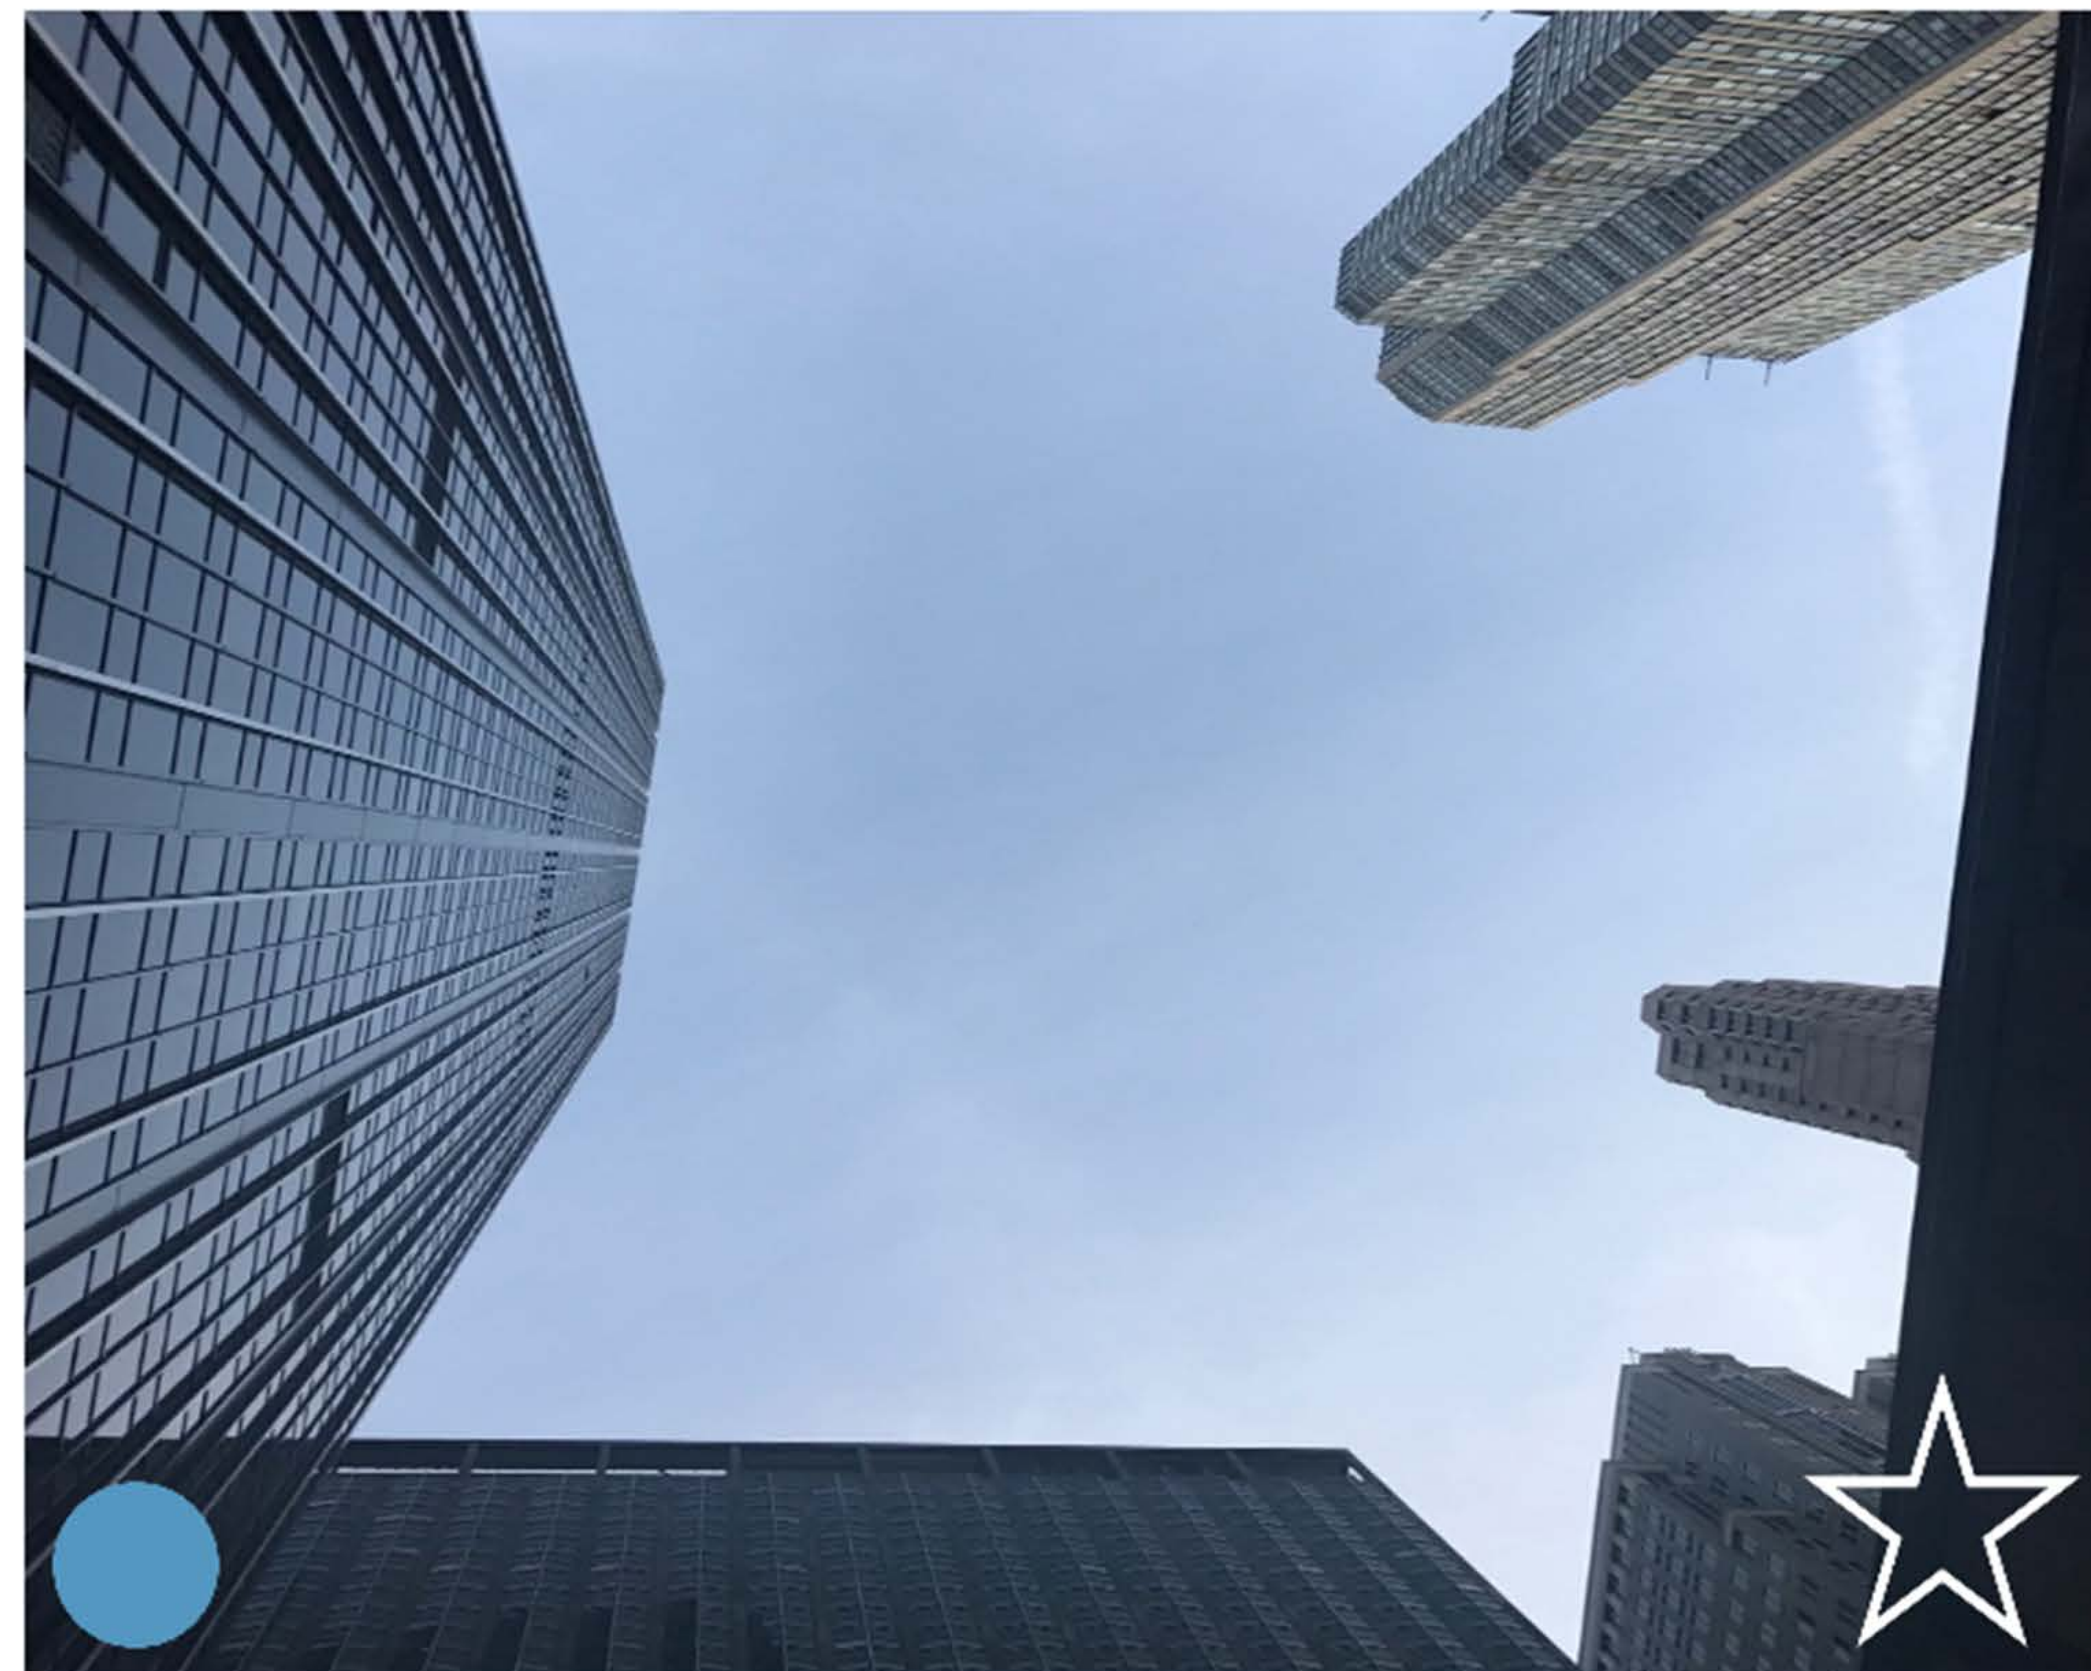
- Filesystem Suggestions** (Instrument)
- Filesystem Activity** (Instrument)
- Disk Usage** (Instrument)
- Disk I/O Latency** (Instrument)



9:41



Add

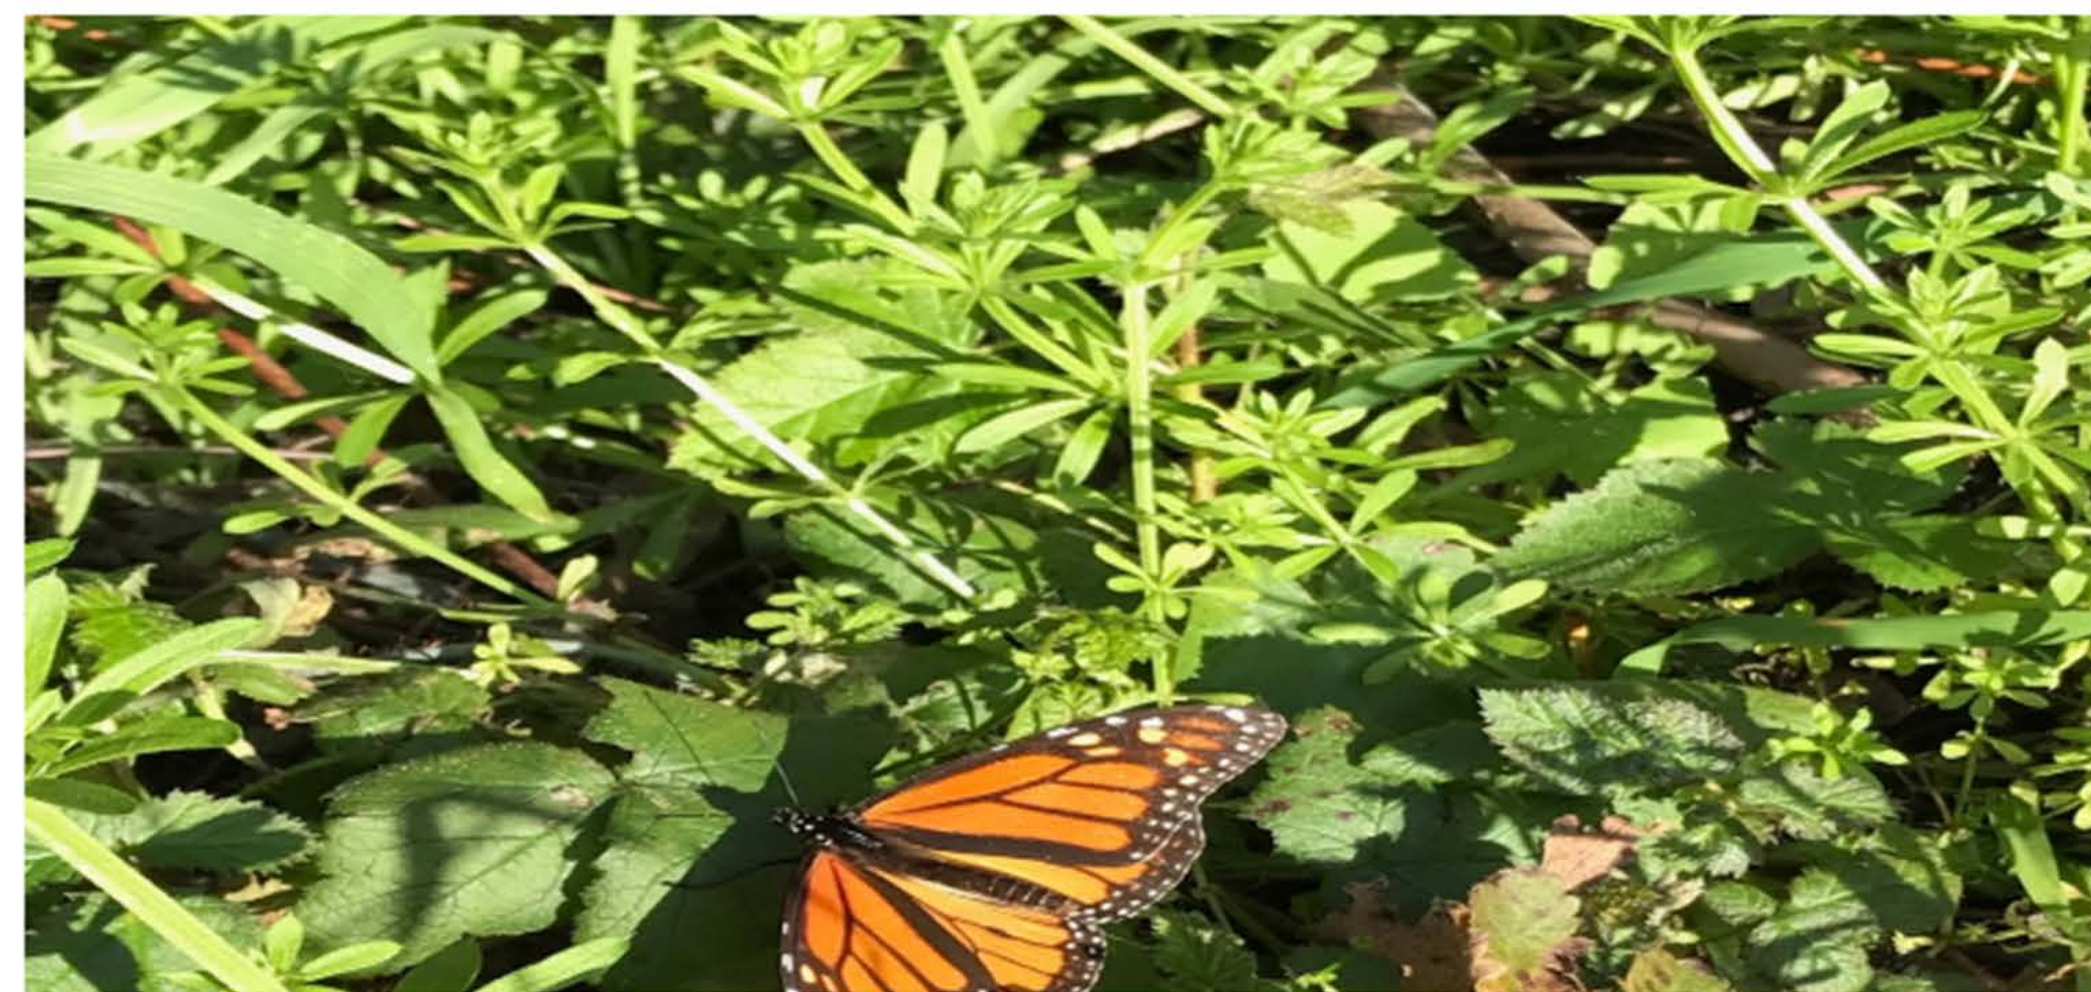
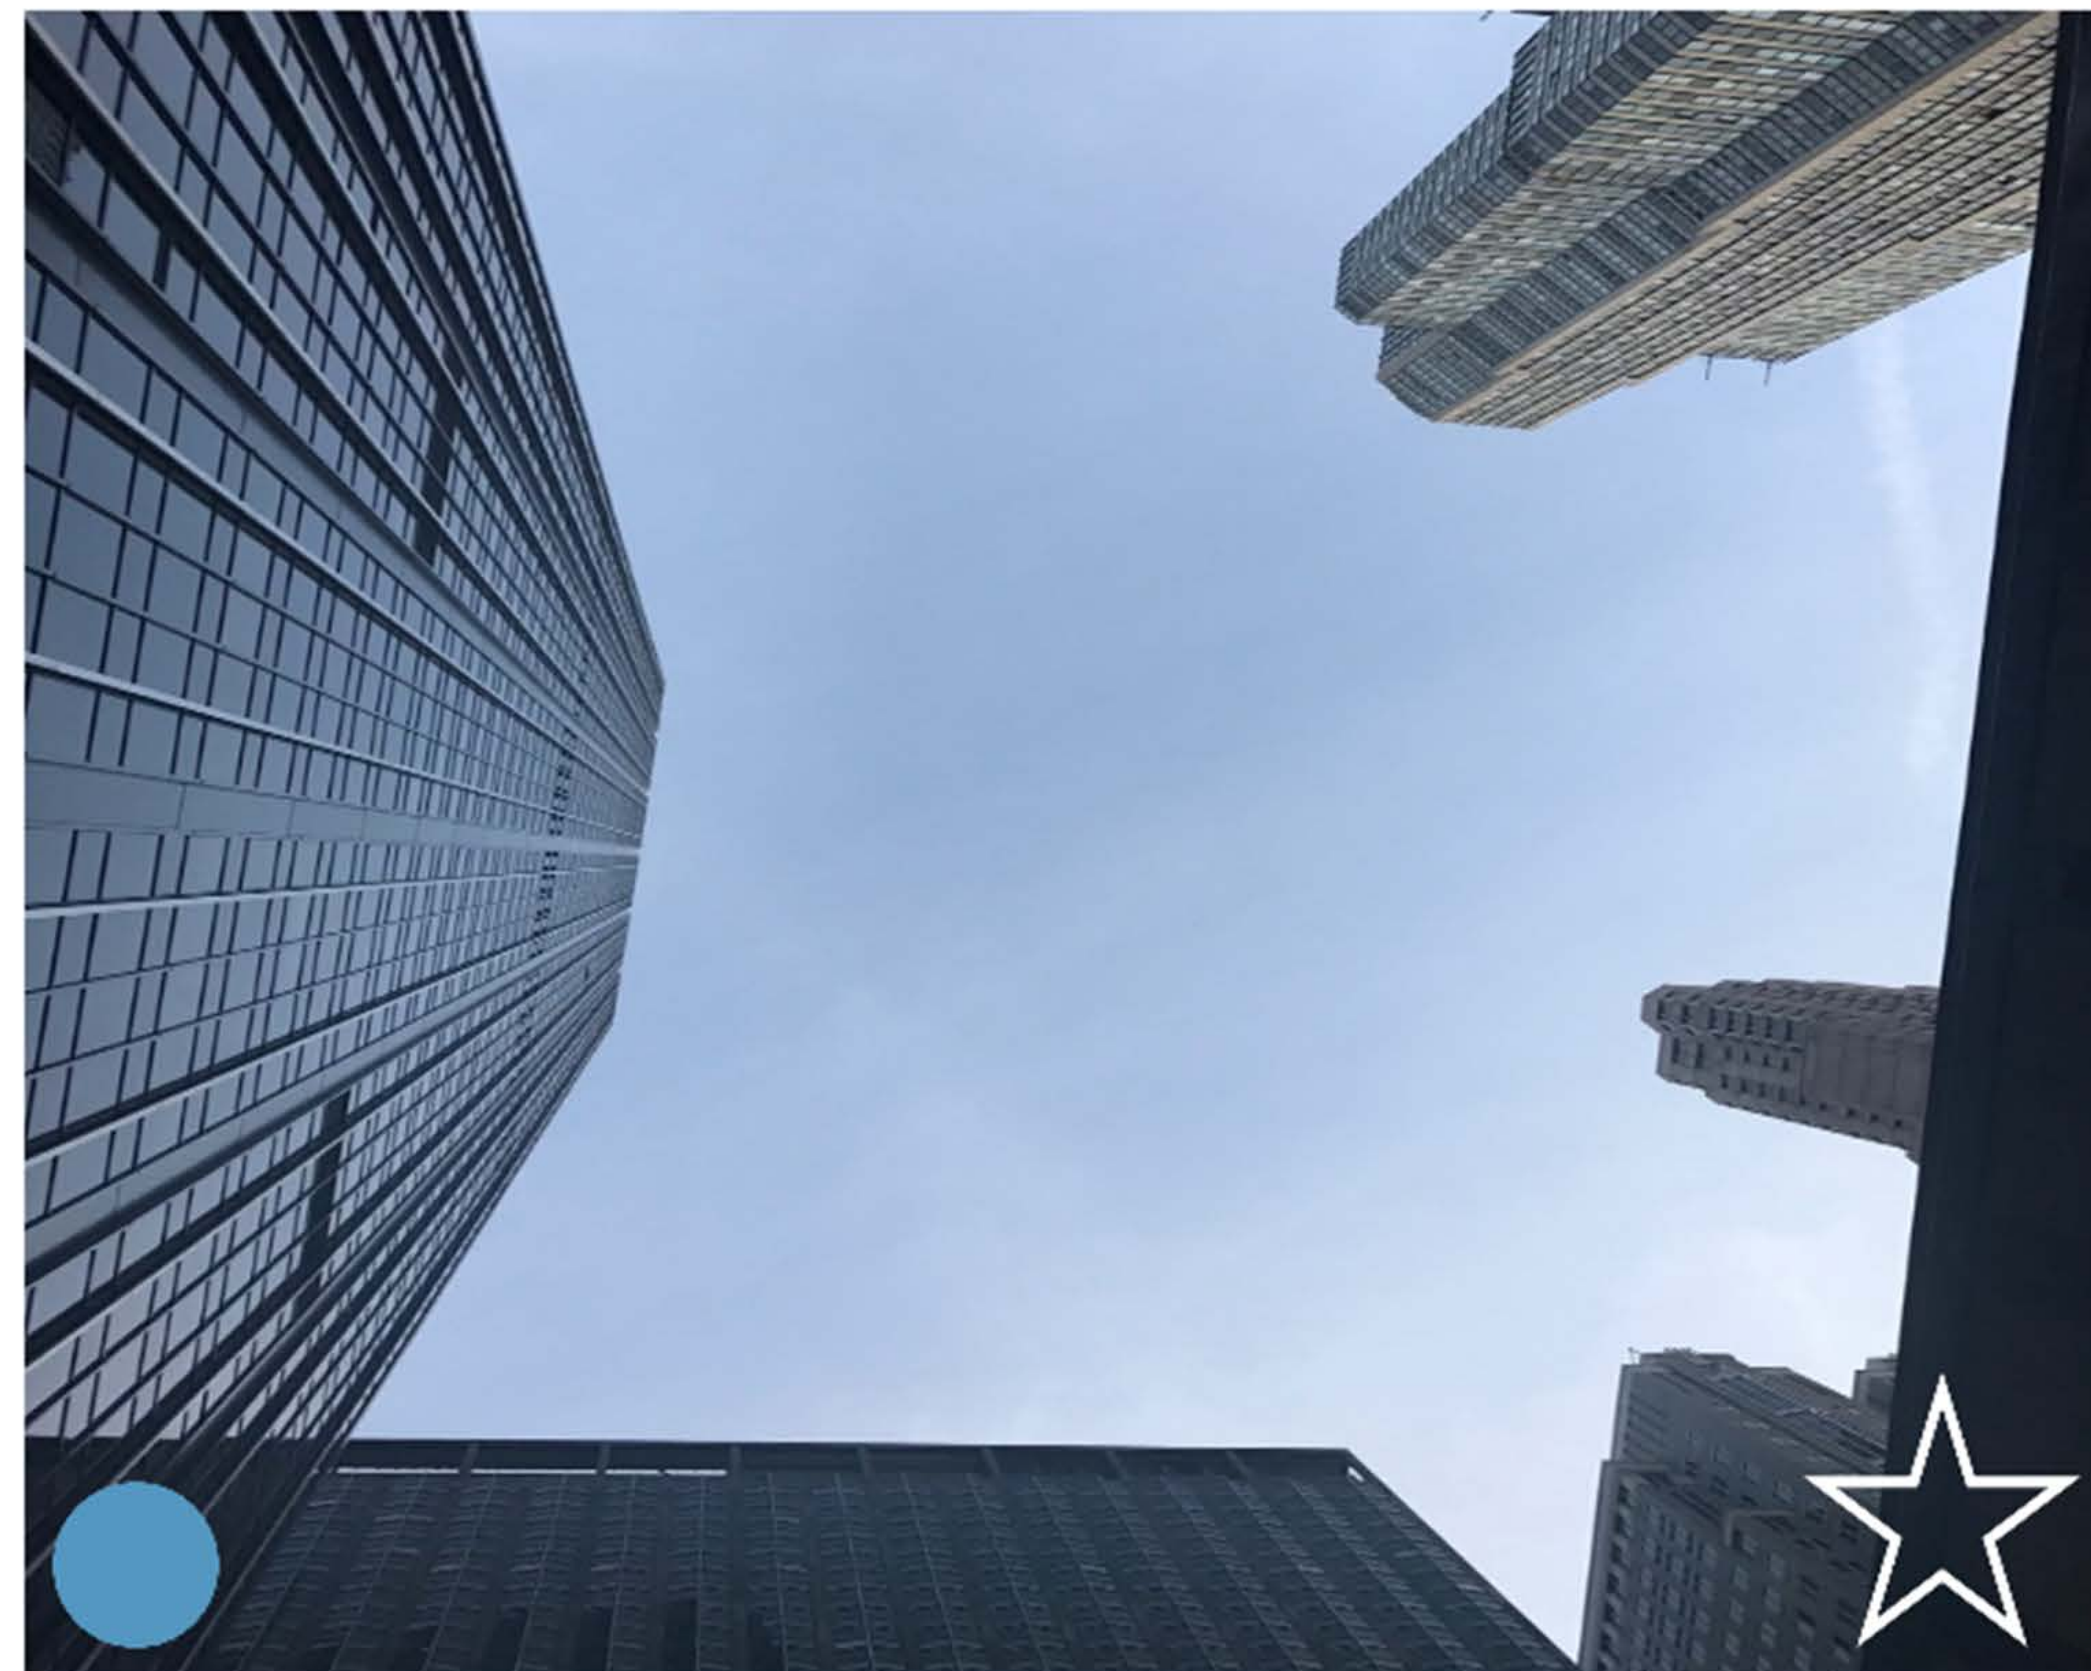




9:41



Add

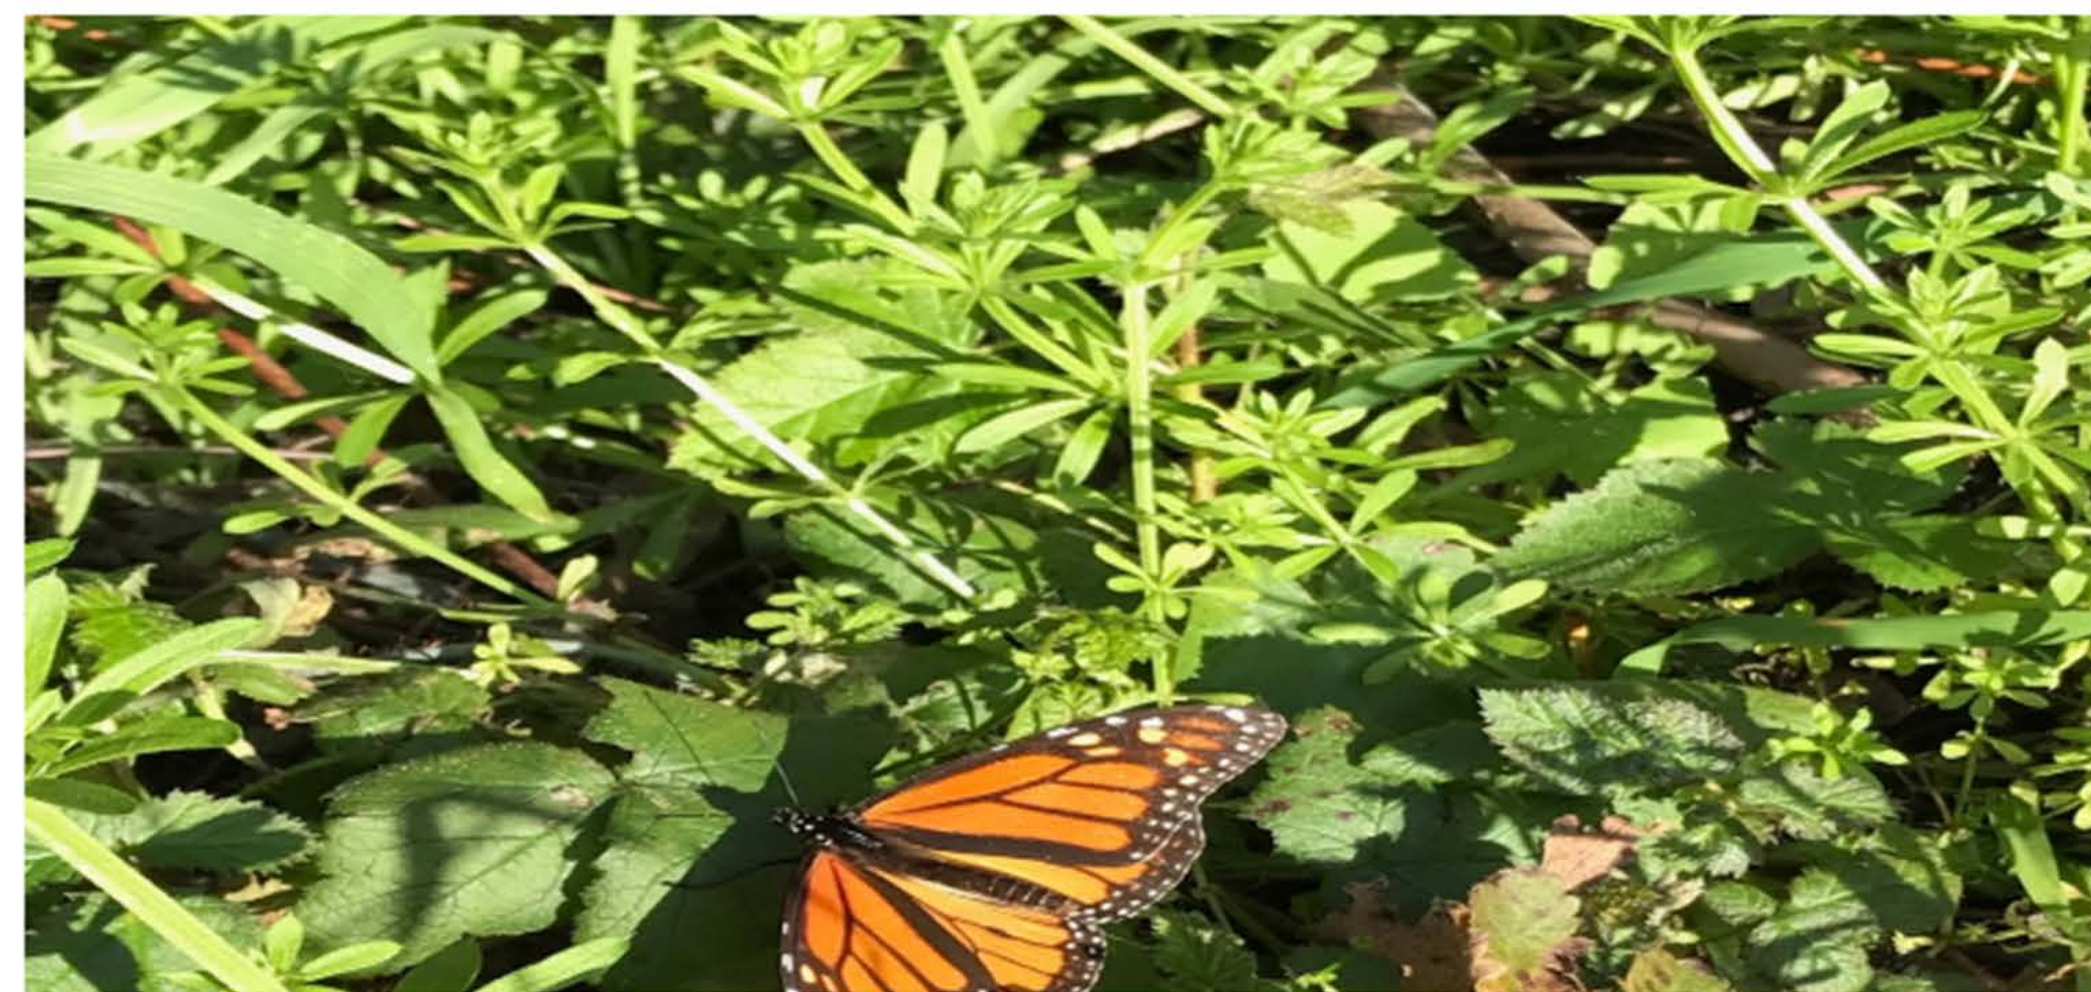




9:41



Add





# Open and Close per Operation

## Disk Usage



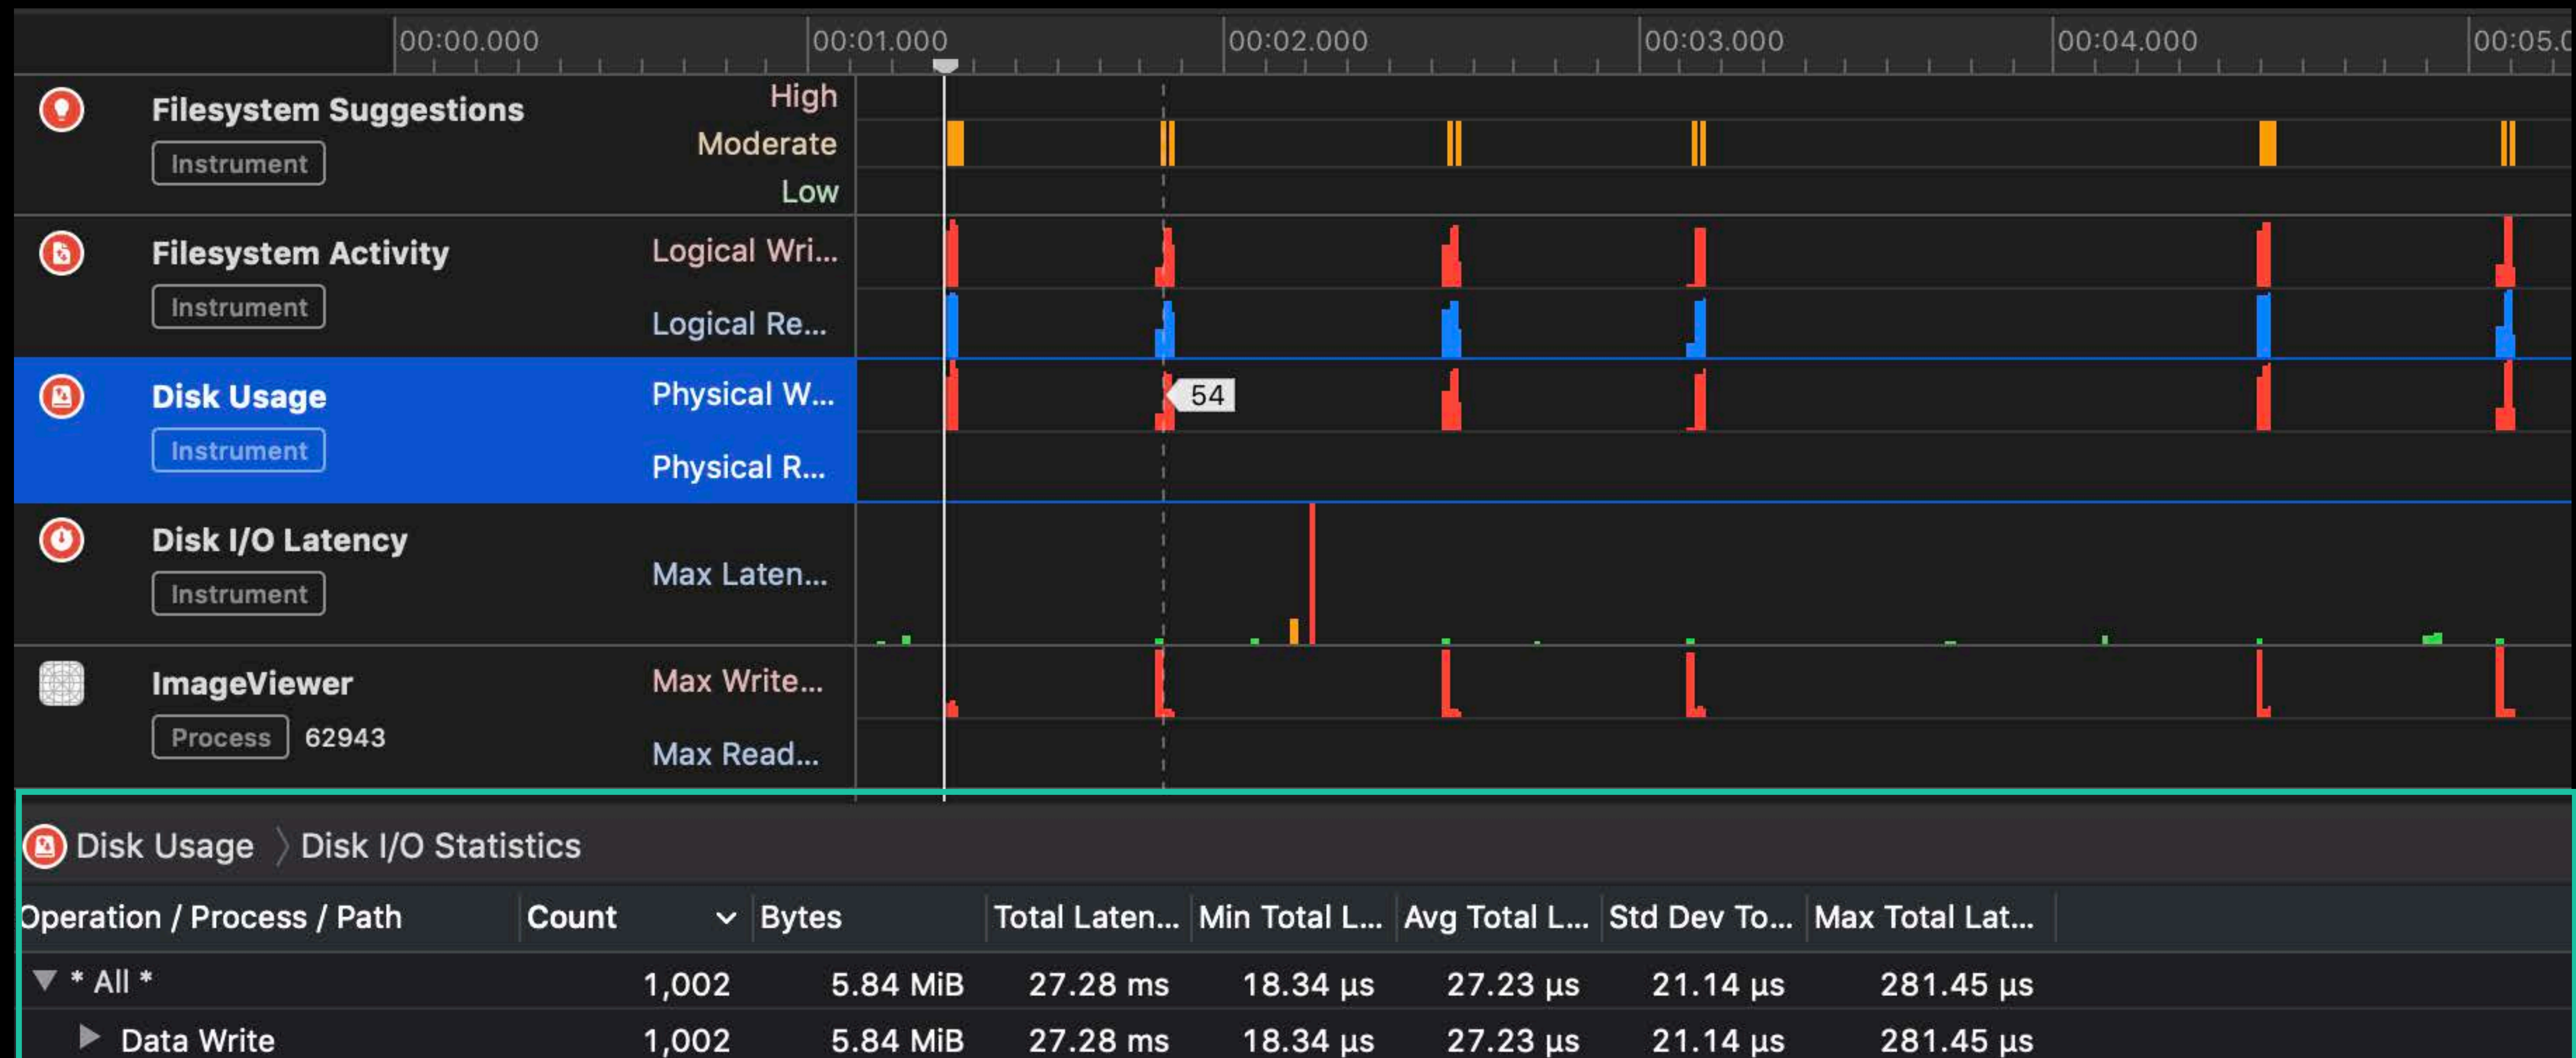
### Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total Lat...
▼ * All *	1,002	5.84 MiB	27.28 ms	18.34 μs	27.23 μs	21.14 μs	281.45 μs
▶ Data Write	1,002	5.84 MiB	27.28 ms	18.34 μs	27.23 μs	21.14 μs	281.45 μs



# Open and Close per Operation

## Disk Usage





# Open and Close per Operation

## Disk Usage — Detail view

Disk Usage > Disk I/O Statistics								
Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...	
▼ * All *	1,003	5.84 MiB	25.71 ms	18.42 µs	25.63 µs	22.17 µs	302.64 µs	
▼ Data Write	1,002	5.84 MiB	25.52 ms	18.42 µs	25.47 µs	21.59 µs	302.64 µs	
▼ ImageViewer (61960)	1,002	5.84 MiB	25.52 ms	18.42 µs	25.47 µs	21.59 µs	302.64 µs	
/ImageViewer/open_close_imageviewer.db-journal	747	3.89 MiB	19.71 ms	18.42 µs	26.39 µs	24.86 µs	302.64 µs	
/ImageViewer/open_close_imageviewer.db	255	1.95 MiB	5.81 ms	19.26 µs	22.79 µs	3.50 µs	39.47 µs	
▶ Data Read	1	4.00 KiB	186.51 µs	186.51 µs	186.51 µs	n/a	186.51 µs	



# Open and Close per Operation

## Disk Usage — Detail view

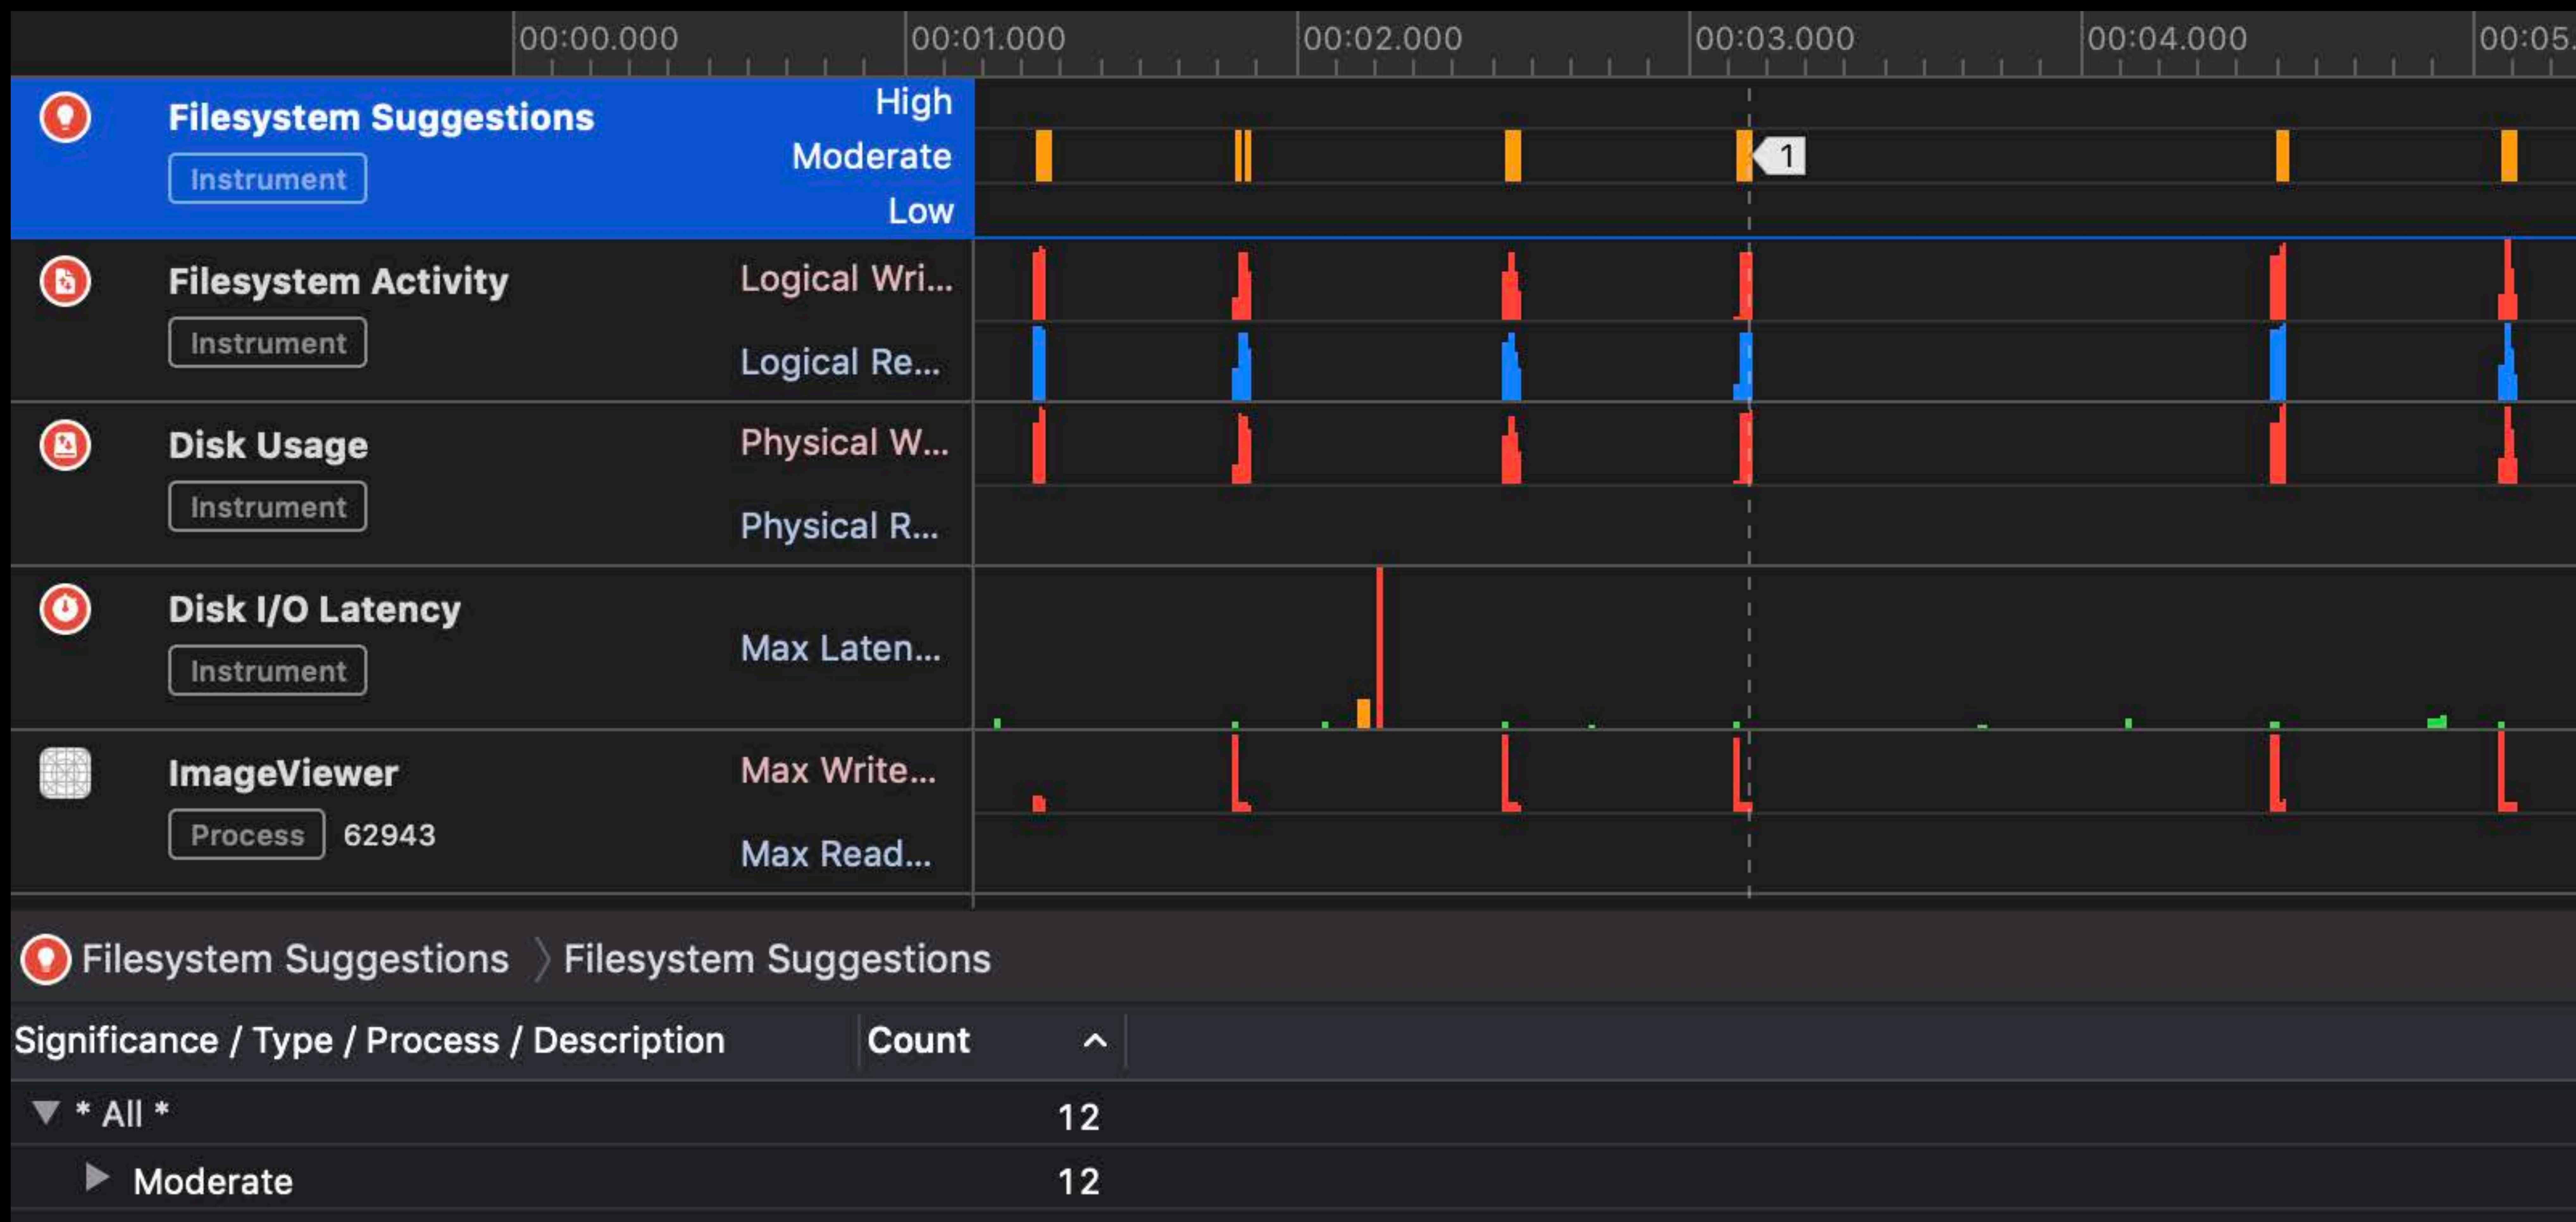
Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	1,003	5.84 MiB	25.71 ms	18.42 µs	25.63 µs	22.17 µs	302.64 µs
▼ Data Write	1,002	5.84 MiB	25.52 ms	18.42 µs	25.47 µs	21.59 µs	302.64 µs
▼ ImageViewer (61960)	1,002	5.84 MiB	25.52 ms	18.42 µs	25.47 µs	21.59 µs	302.64 µs
/ImageViewer/open_close_imageviewer.db-journal	747	3.89 MiB	19.71 ms	18.42 µs	26.39 µs	24.86 µs	302.64 µs
/ImageViewer/open_close_imageviewer.db	255	1.95 MiB	5.81 ms	19.26 µs	22.79 µs	3.50 µs	39.47 µs
▶ Data Read	1	4.00 KiB	186.51 µs	186.51 µs	186.51 µs	n/a	186.51 µs



# Open and Close per Operation

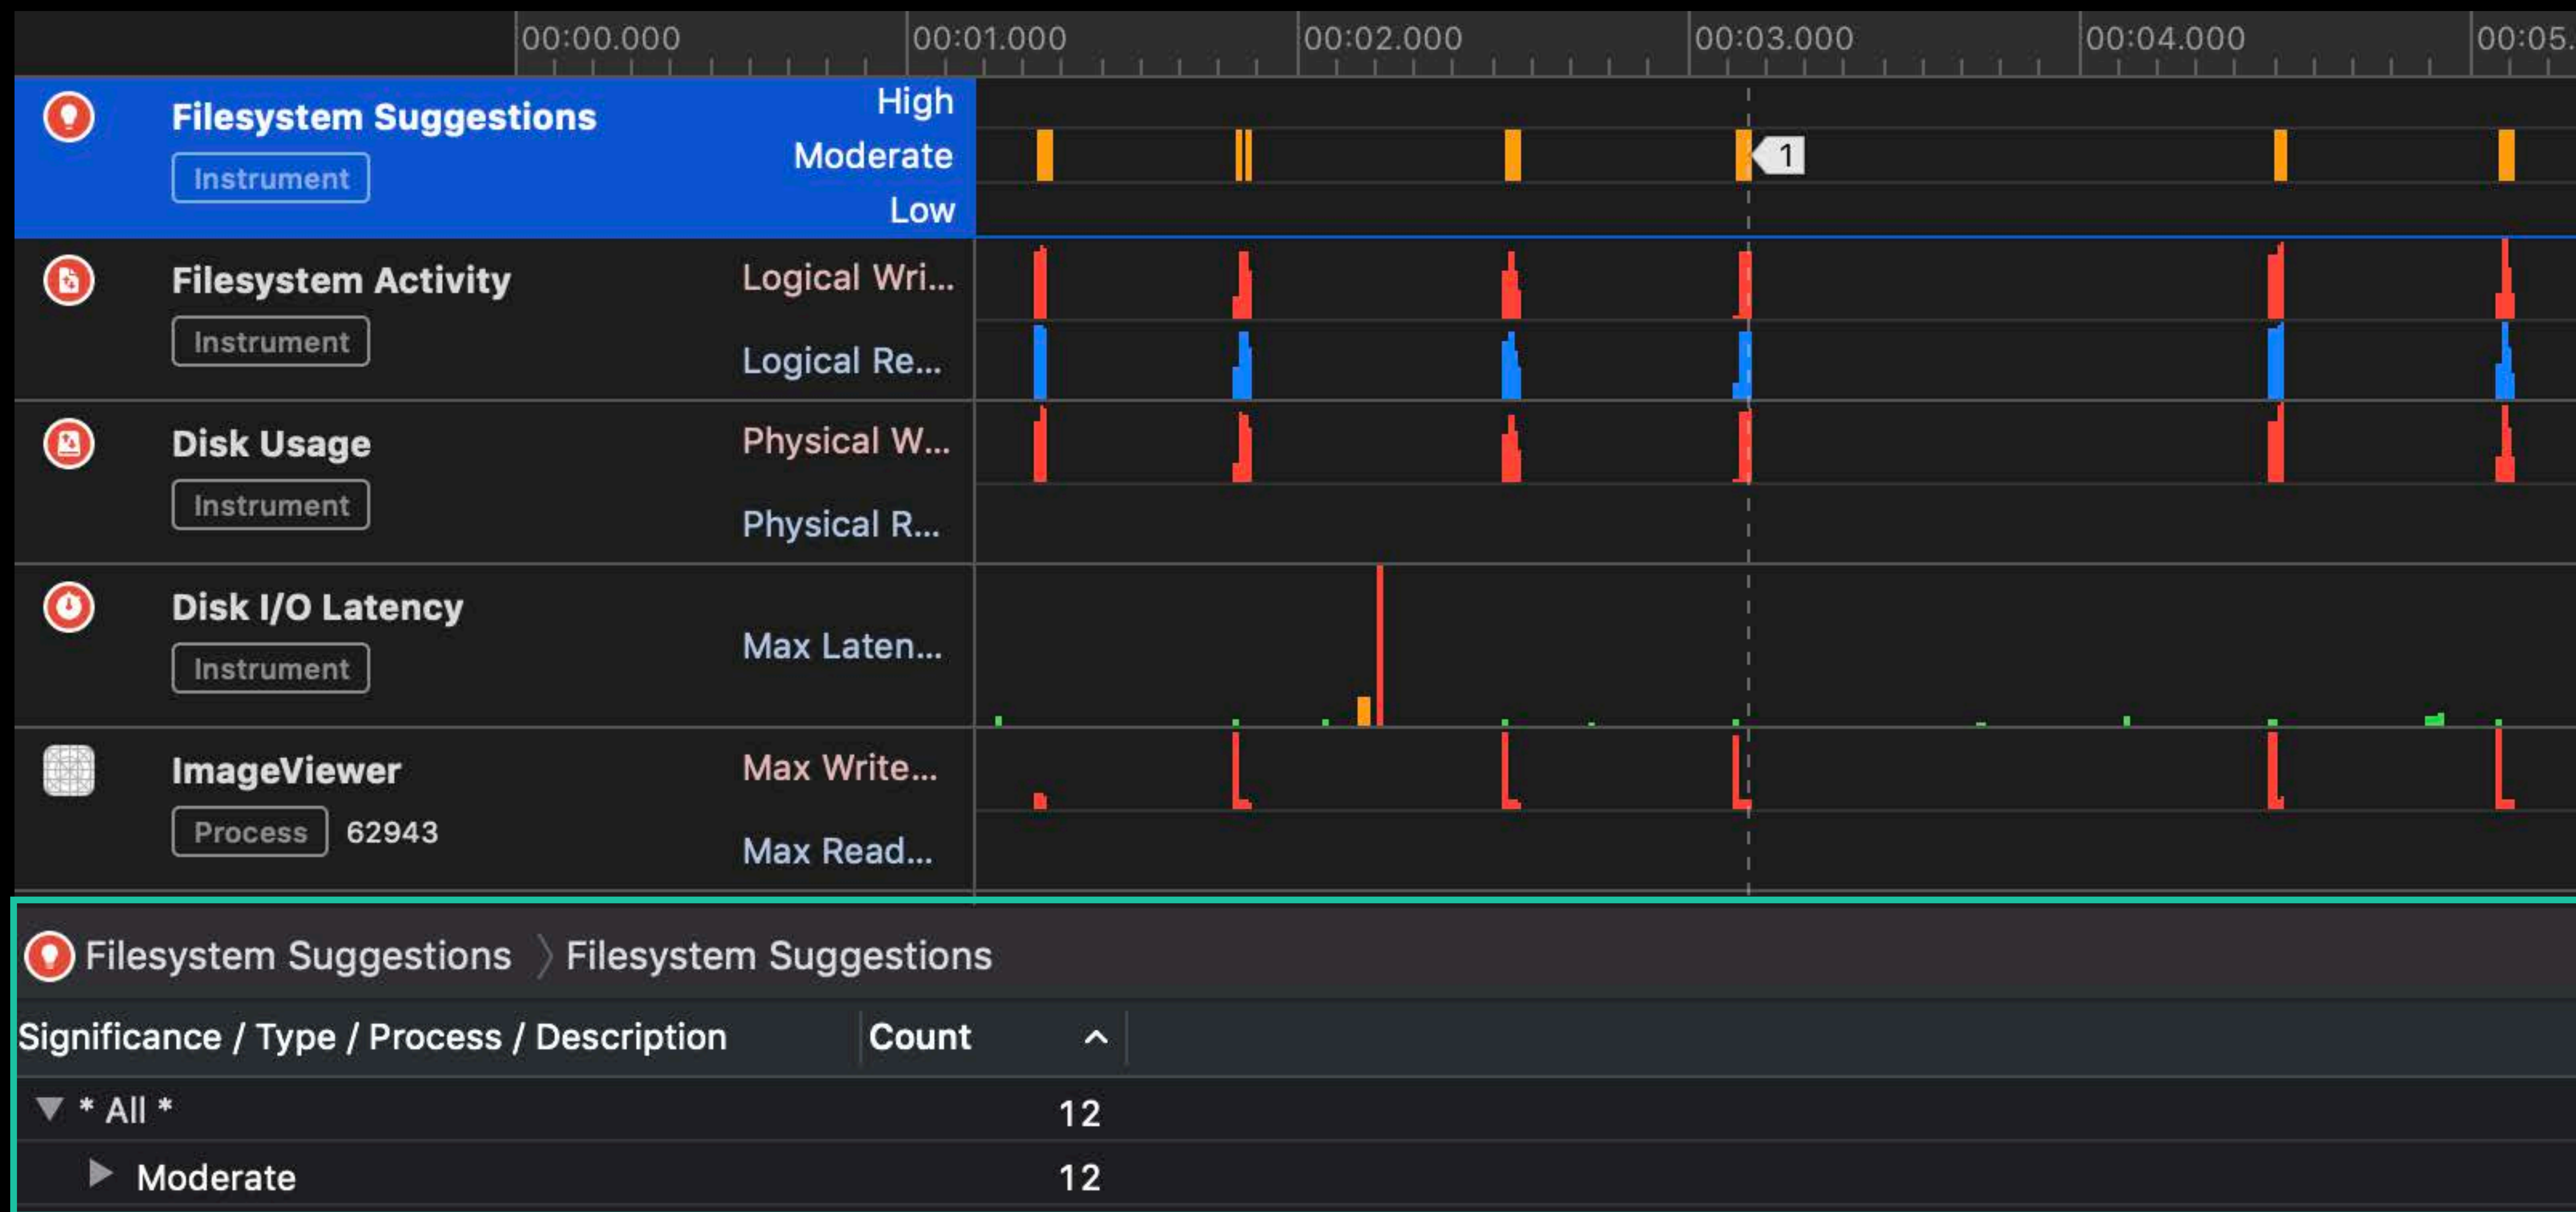
## Filesystem Suggestions





# Open and Close per Operation

## Filesystem Suggestions





# Open and Close per Operation

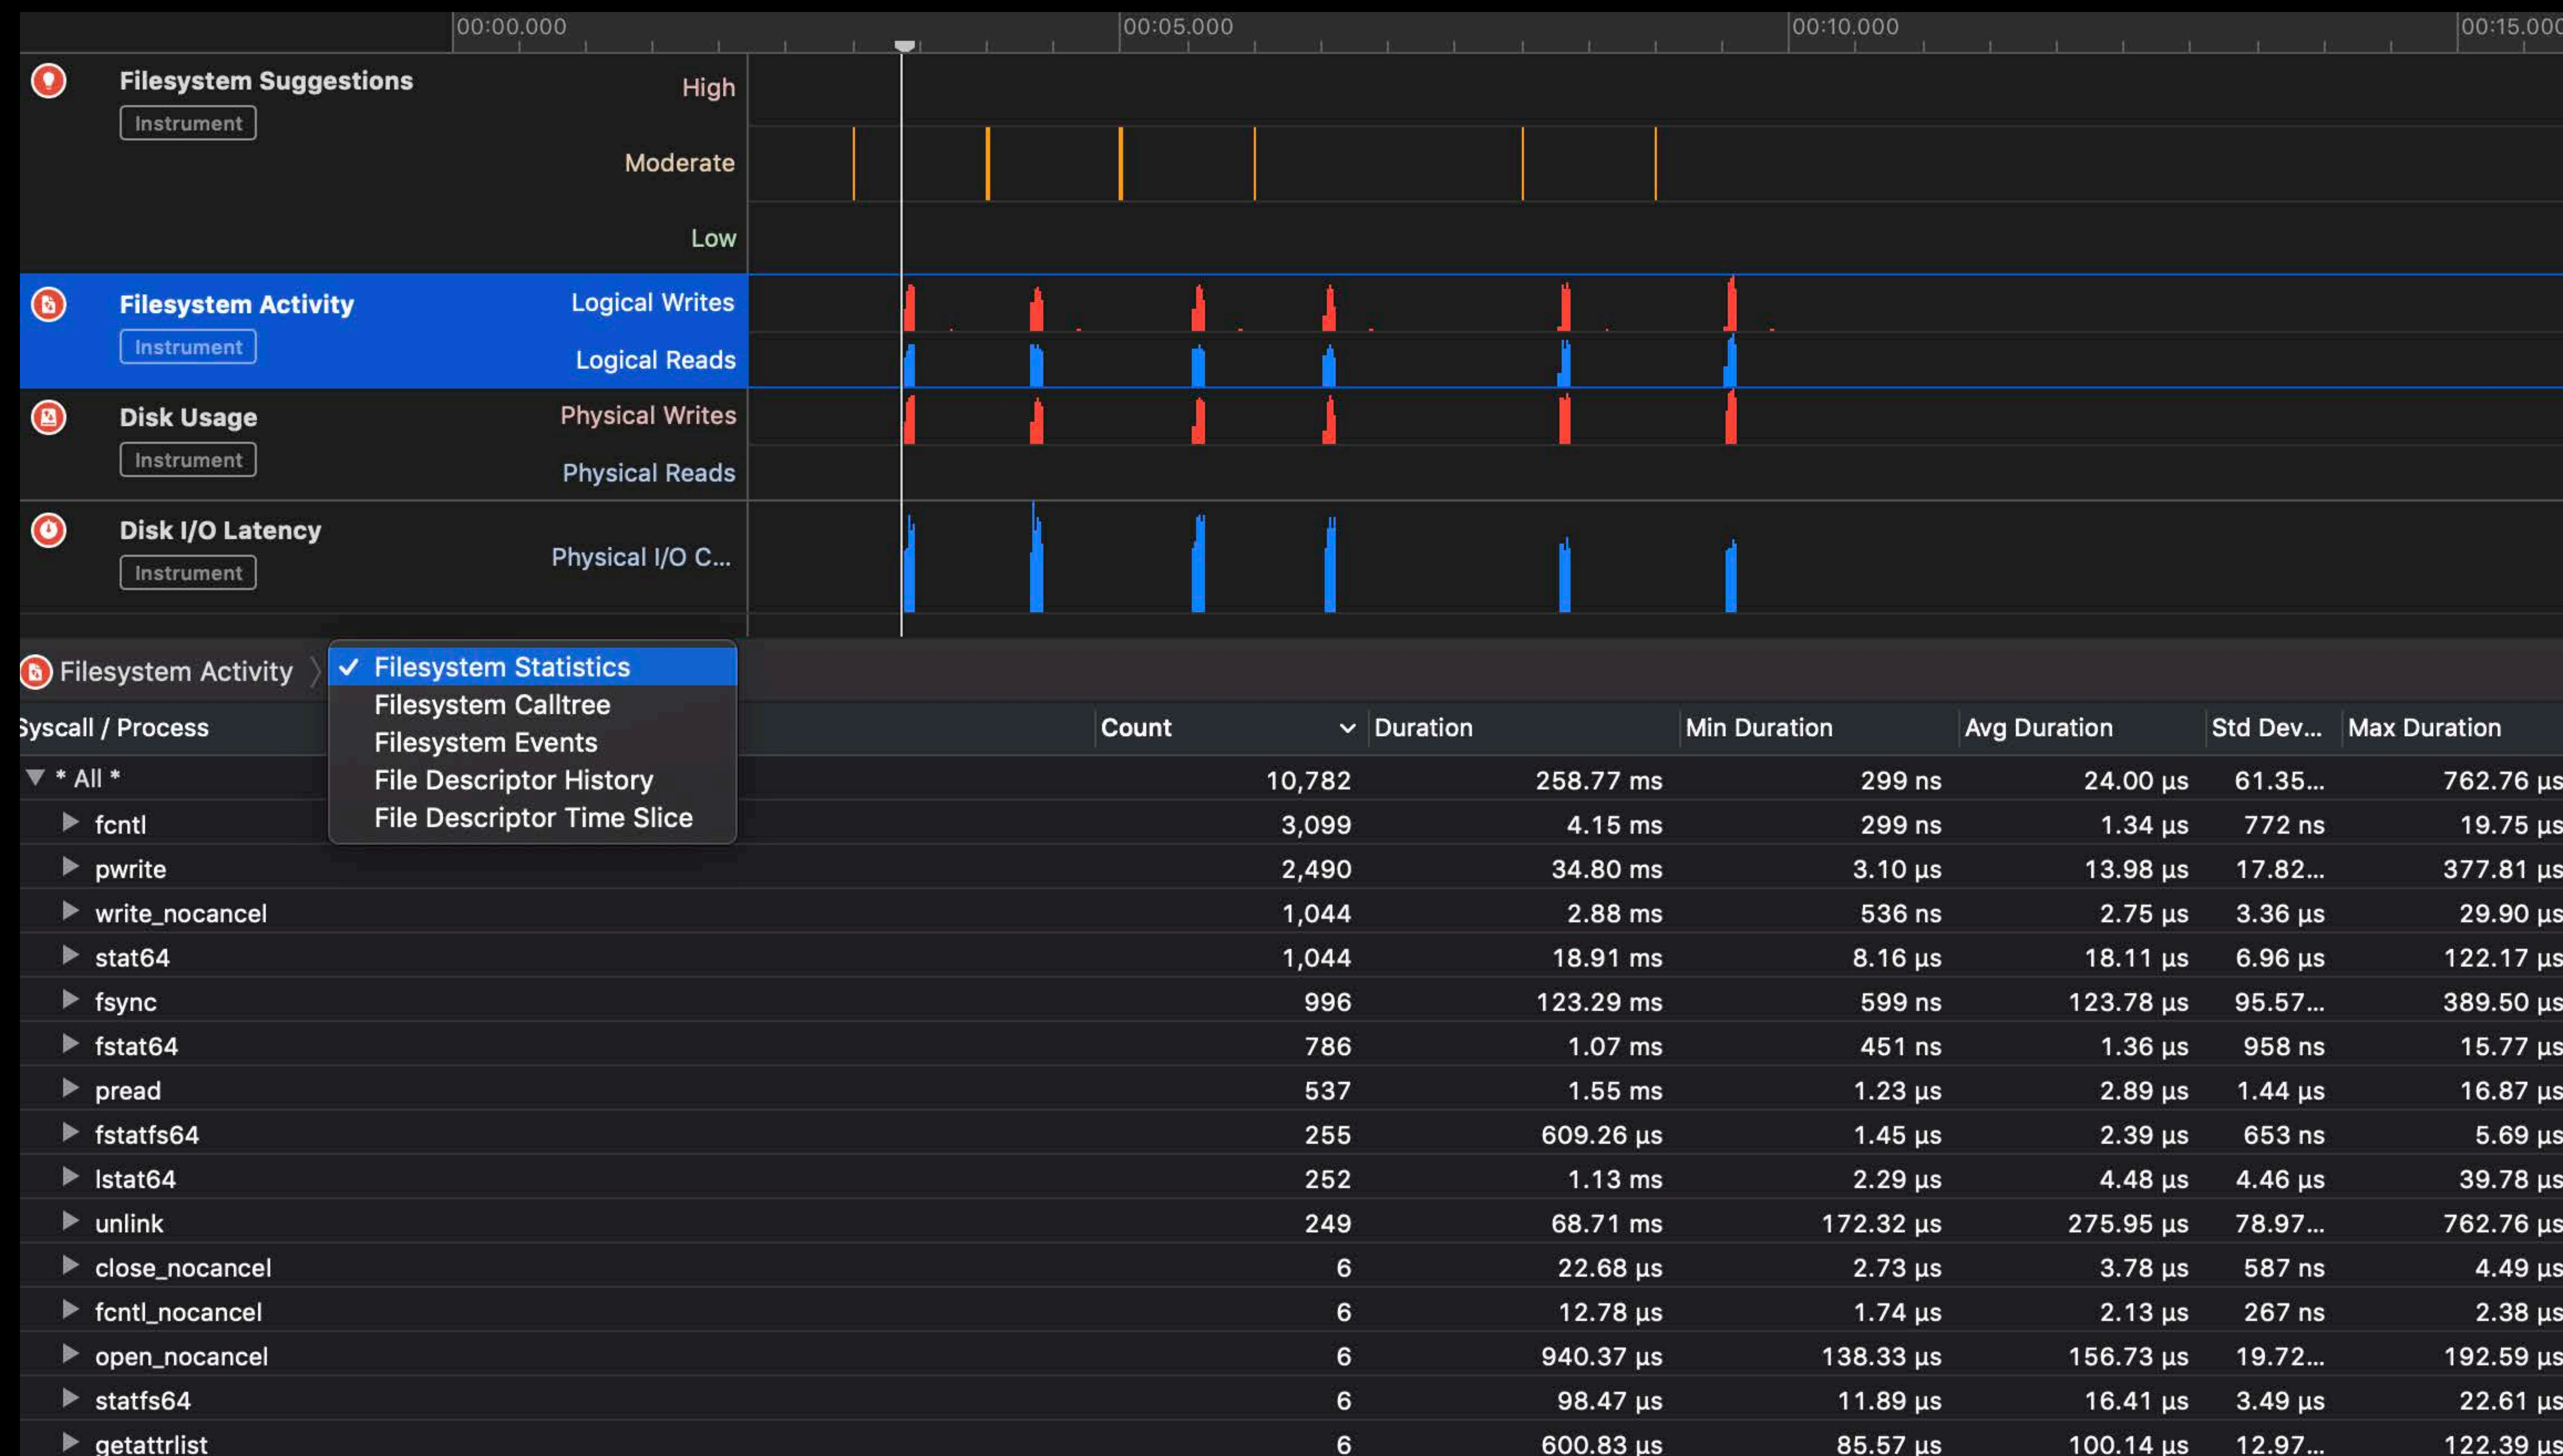
## Filesystem Suggestions — Detail view

Filesystem Suggestions > Filesystem Suggestions	
Significance / Type / Process / Description	Count
▼ * All *	12
▼ Moderate	12
▼ Excessive Writes	12
▼ ImageViewer (61960)	12
ImageViewer (61960) performed 83 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 84 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 81 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 82 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 79 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 88 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 99 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 64 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 77 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 90 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 102 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1
ImageViewer (61960) performed 61 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	1



# Open and Close per Operation

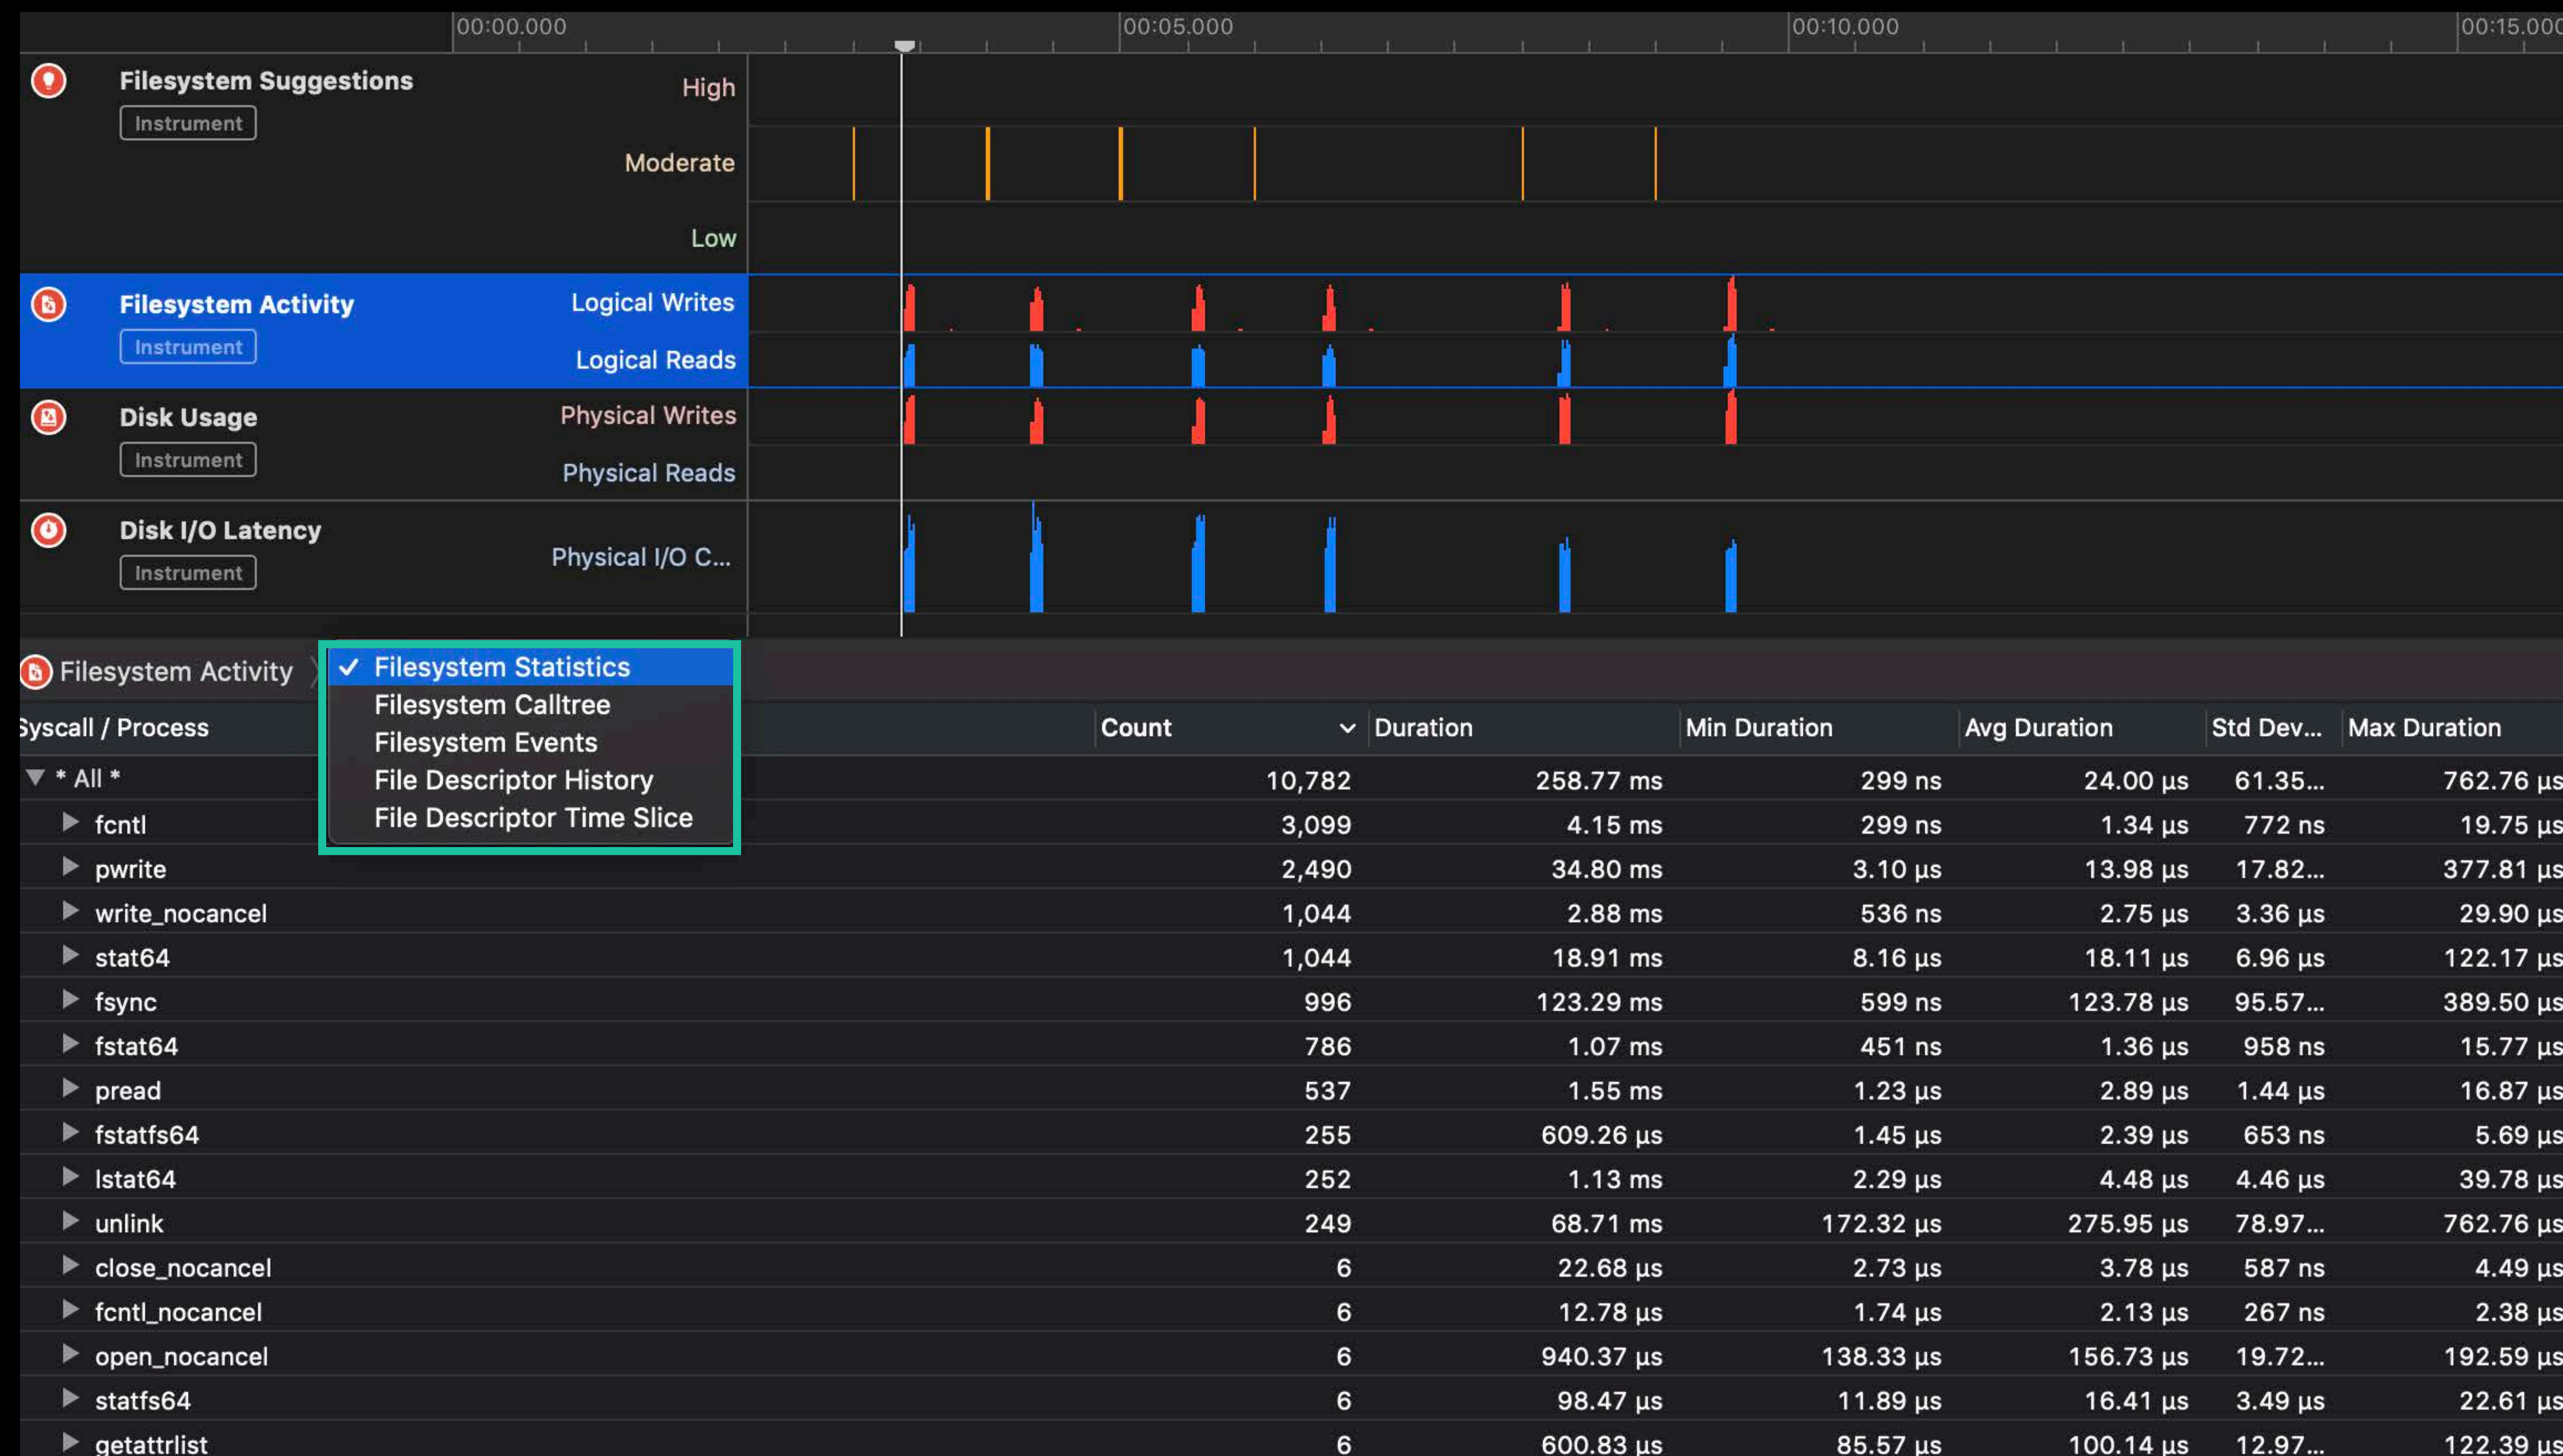
## Filesystem Activity





# Open and Close per Operation

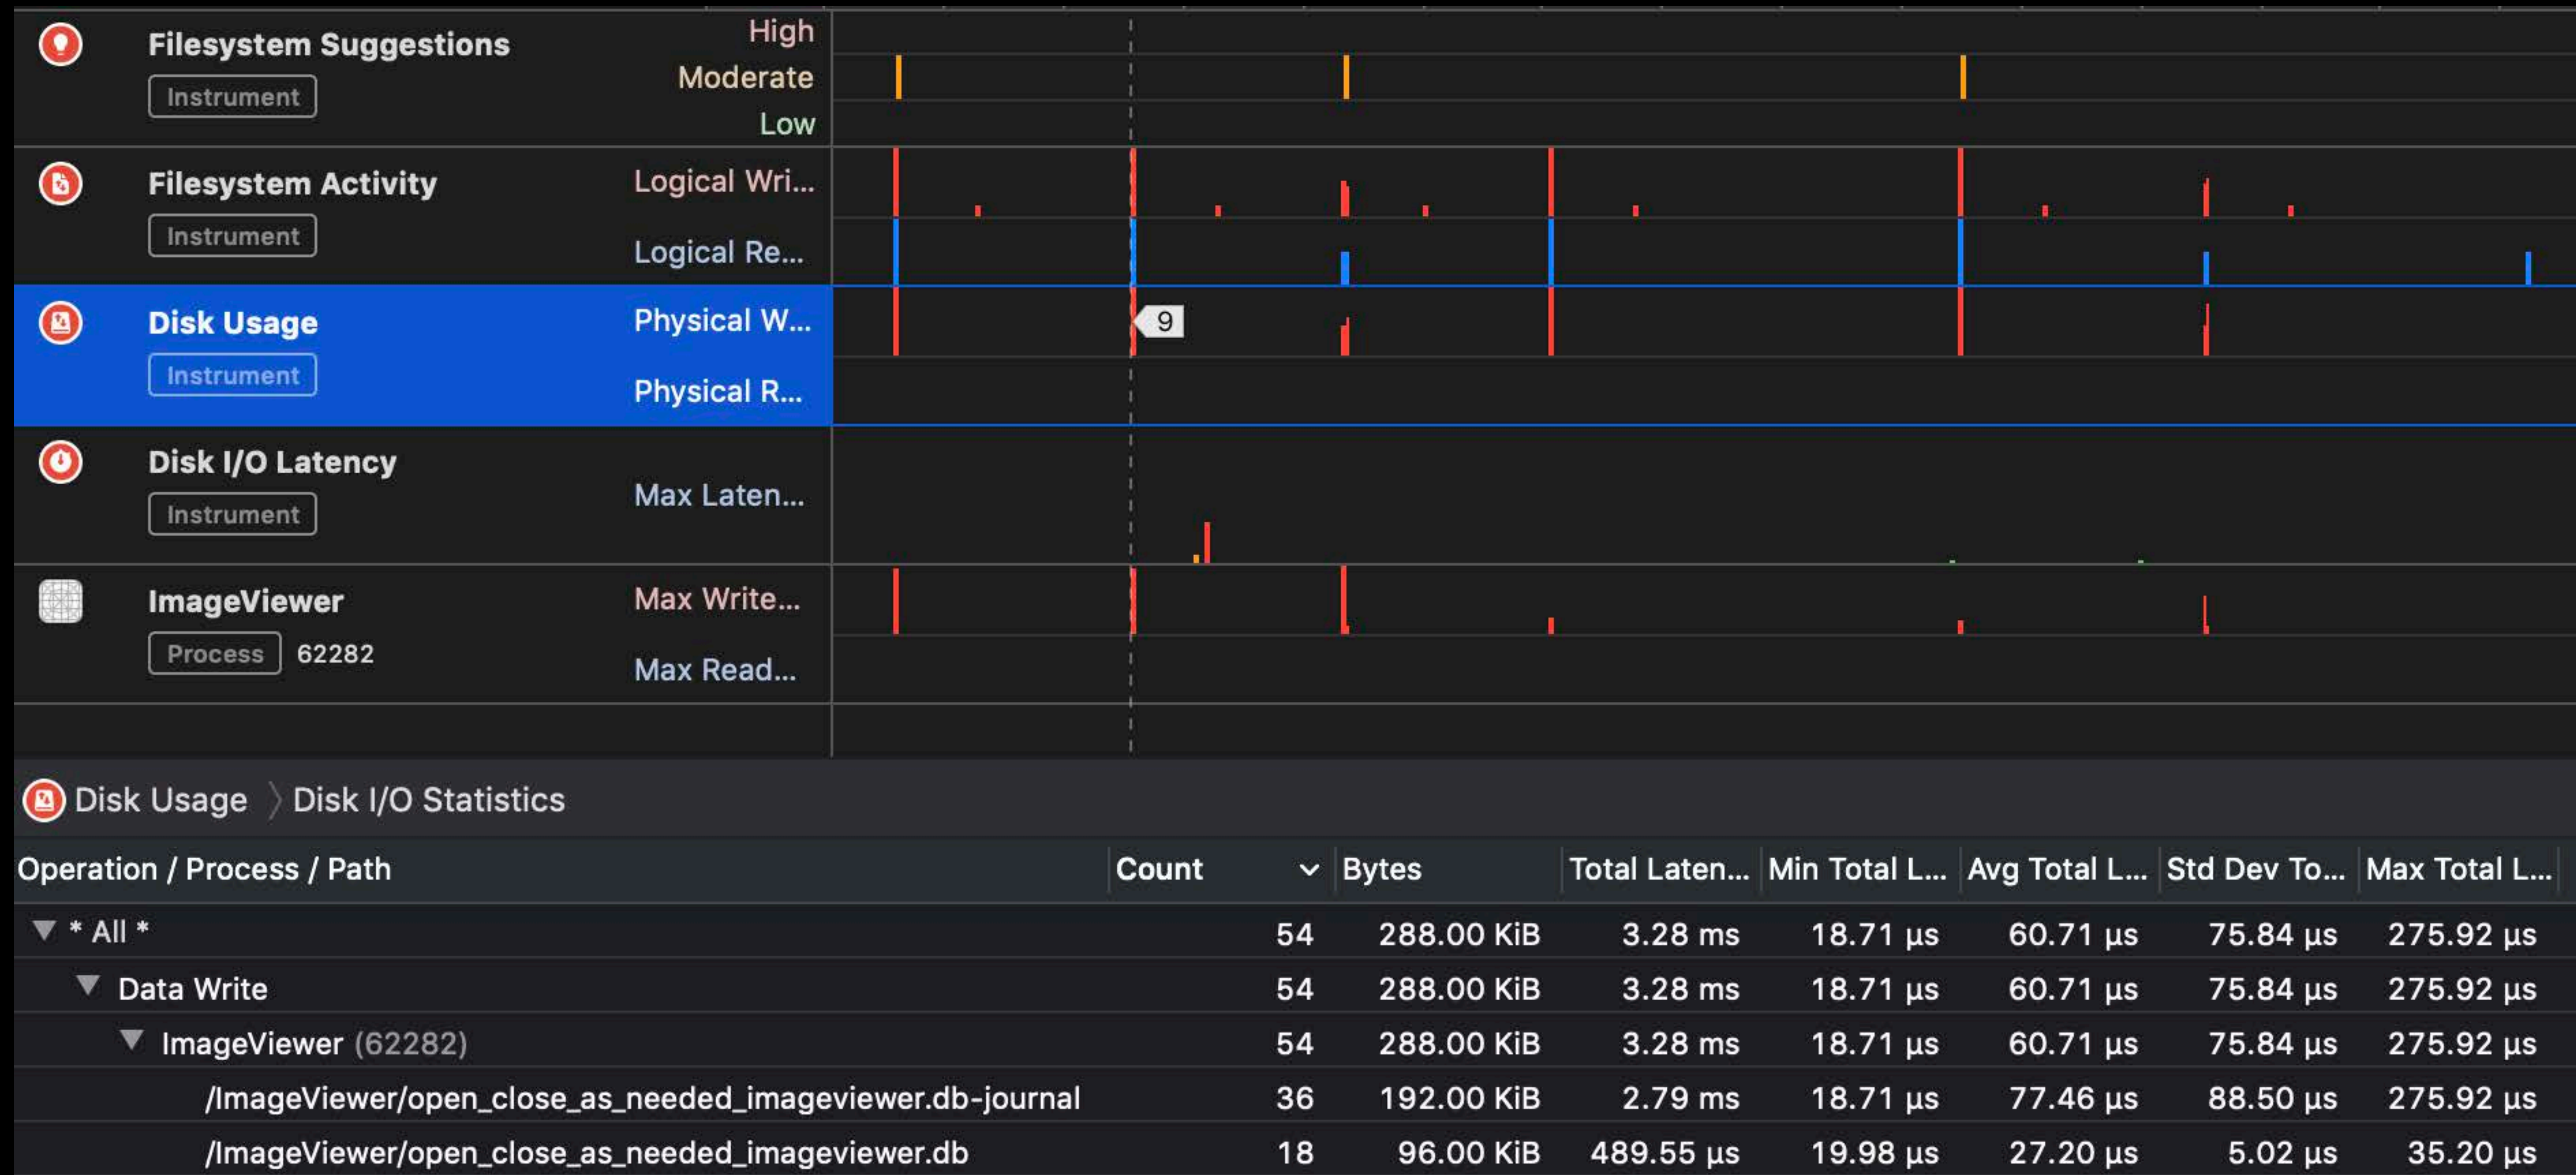
## Filesystem Activity





# Open and Close as Needed

## Disk Usage





# Open and Close as Needed

## Disk Usage — Detail view


 Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	54	288.00 KiB	3.28 ms	18.71 µs	60.71 µs	75.84 µs	275.92 µs
▼ Data Write	54	288.00 KiB	3.28 ms	18.71 µs	60.71 µs	75.84 µs	275.92 µs
▼ ImageViewer (62282)	54	288.00 KiB	3.28 ms	18.71 µs	60.71 µs	75.84 µs	275.92 µs
/ImageViewer/open_close_as_needed_imageviewer.db-journal	36	192.00 KiB	2.79 ms	18.71 µs	77.46 µs	88.50 µs	275.92 µs
/ImageViewer/open_close_as_needed_imageviewer.db	18	96.00 KiB	489.55 µs	19.98 µs	27.20 µs	5.02 µs	35.20 µs



# Open and Close as Needed

## Disk Usage — Detail view

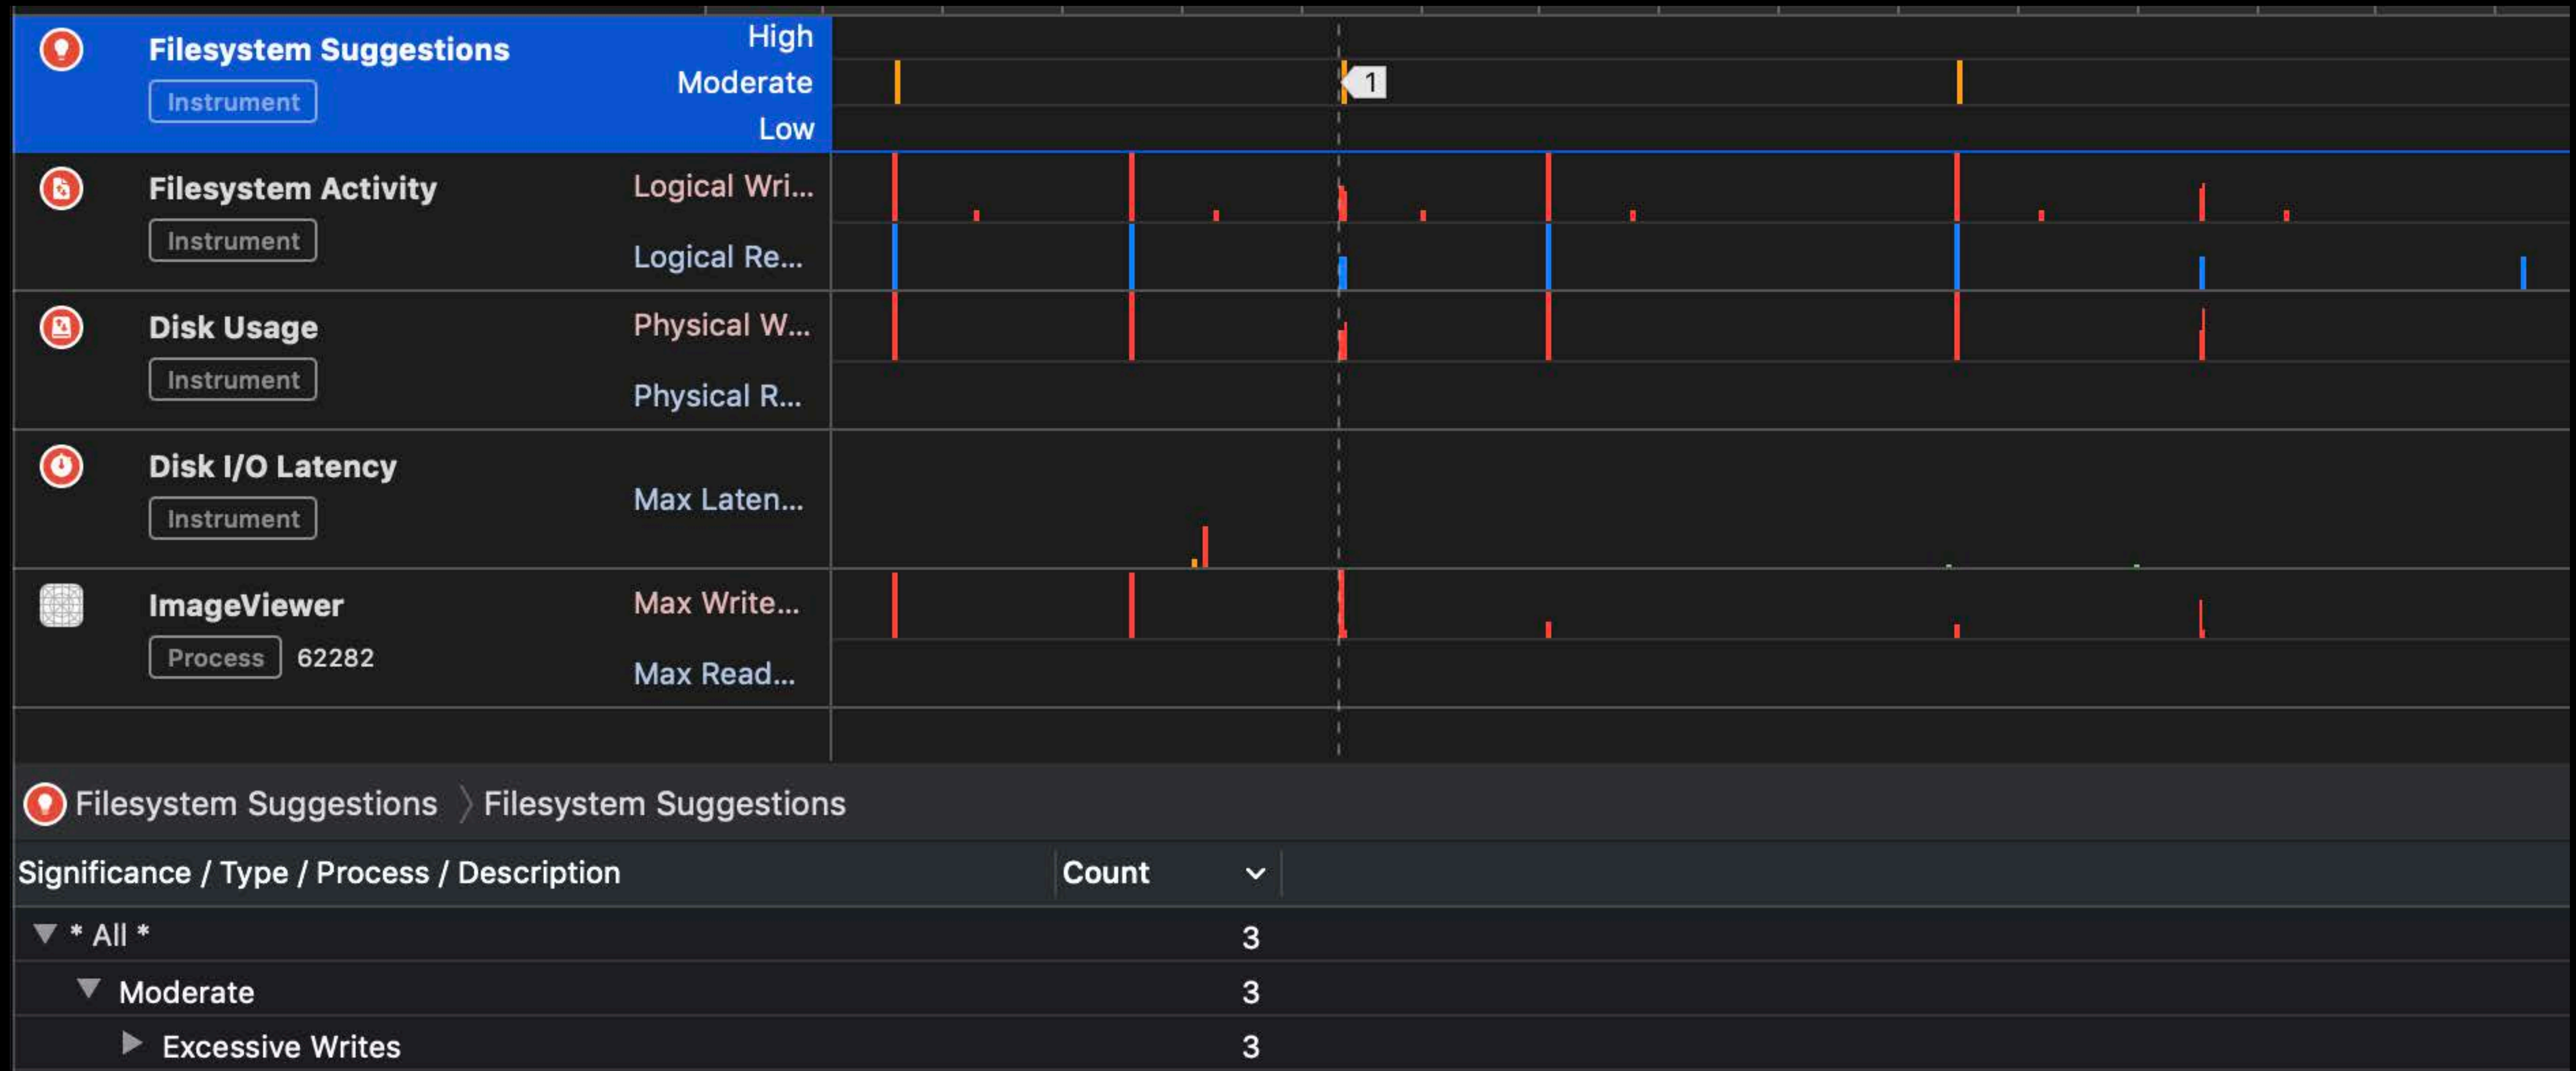
 Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	54	288.00 KiB	3.28 ms	18.71 μs	60.71 μs	75.84 μs	275.92 μs
▼ Data Write	54	288.00 KiB	3.28 ms	18.71 μs	60.71 μs	75.84 μs	275.92 μs
▼ ImageViewer (62282)	54	288.00 KiB	3.28 ms	18.71 μs	60.71 μs	75.84 μs	275.92 μs
/ImageViewer/open_close_as_needed_imageviewer.db-journal	36	192.00 KiB	2.79 ms	18.71 μs	77.46 μs	88.50 μs	275.92 μs
/ImageViewer/open_close_as_needed_imageviewer.db	18	96.00 KiB	489.55 μs	19.98 μs	27.20 μs	5.02 μs	35.20 μs



# Open and Close as Needed

## Filesystem Suggestions





# Open and Close as Needed


## Filesystem Suggestions — Detail view

Filesystem Suggestions > Filesystem Suggestions	
Significance / Type / Process / Description	Count
▼ * All *	3
▼ Moderate	3
▼ Excessive Writes	3
▼ ImageViewer (62282)	3
ImageViewer (62282) performed 17 physical writes within 1 second. This can put excessive strain on physical storage media and hinder application performance.	3



# Delete Mode Journaling

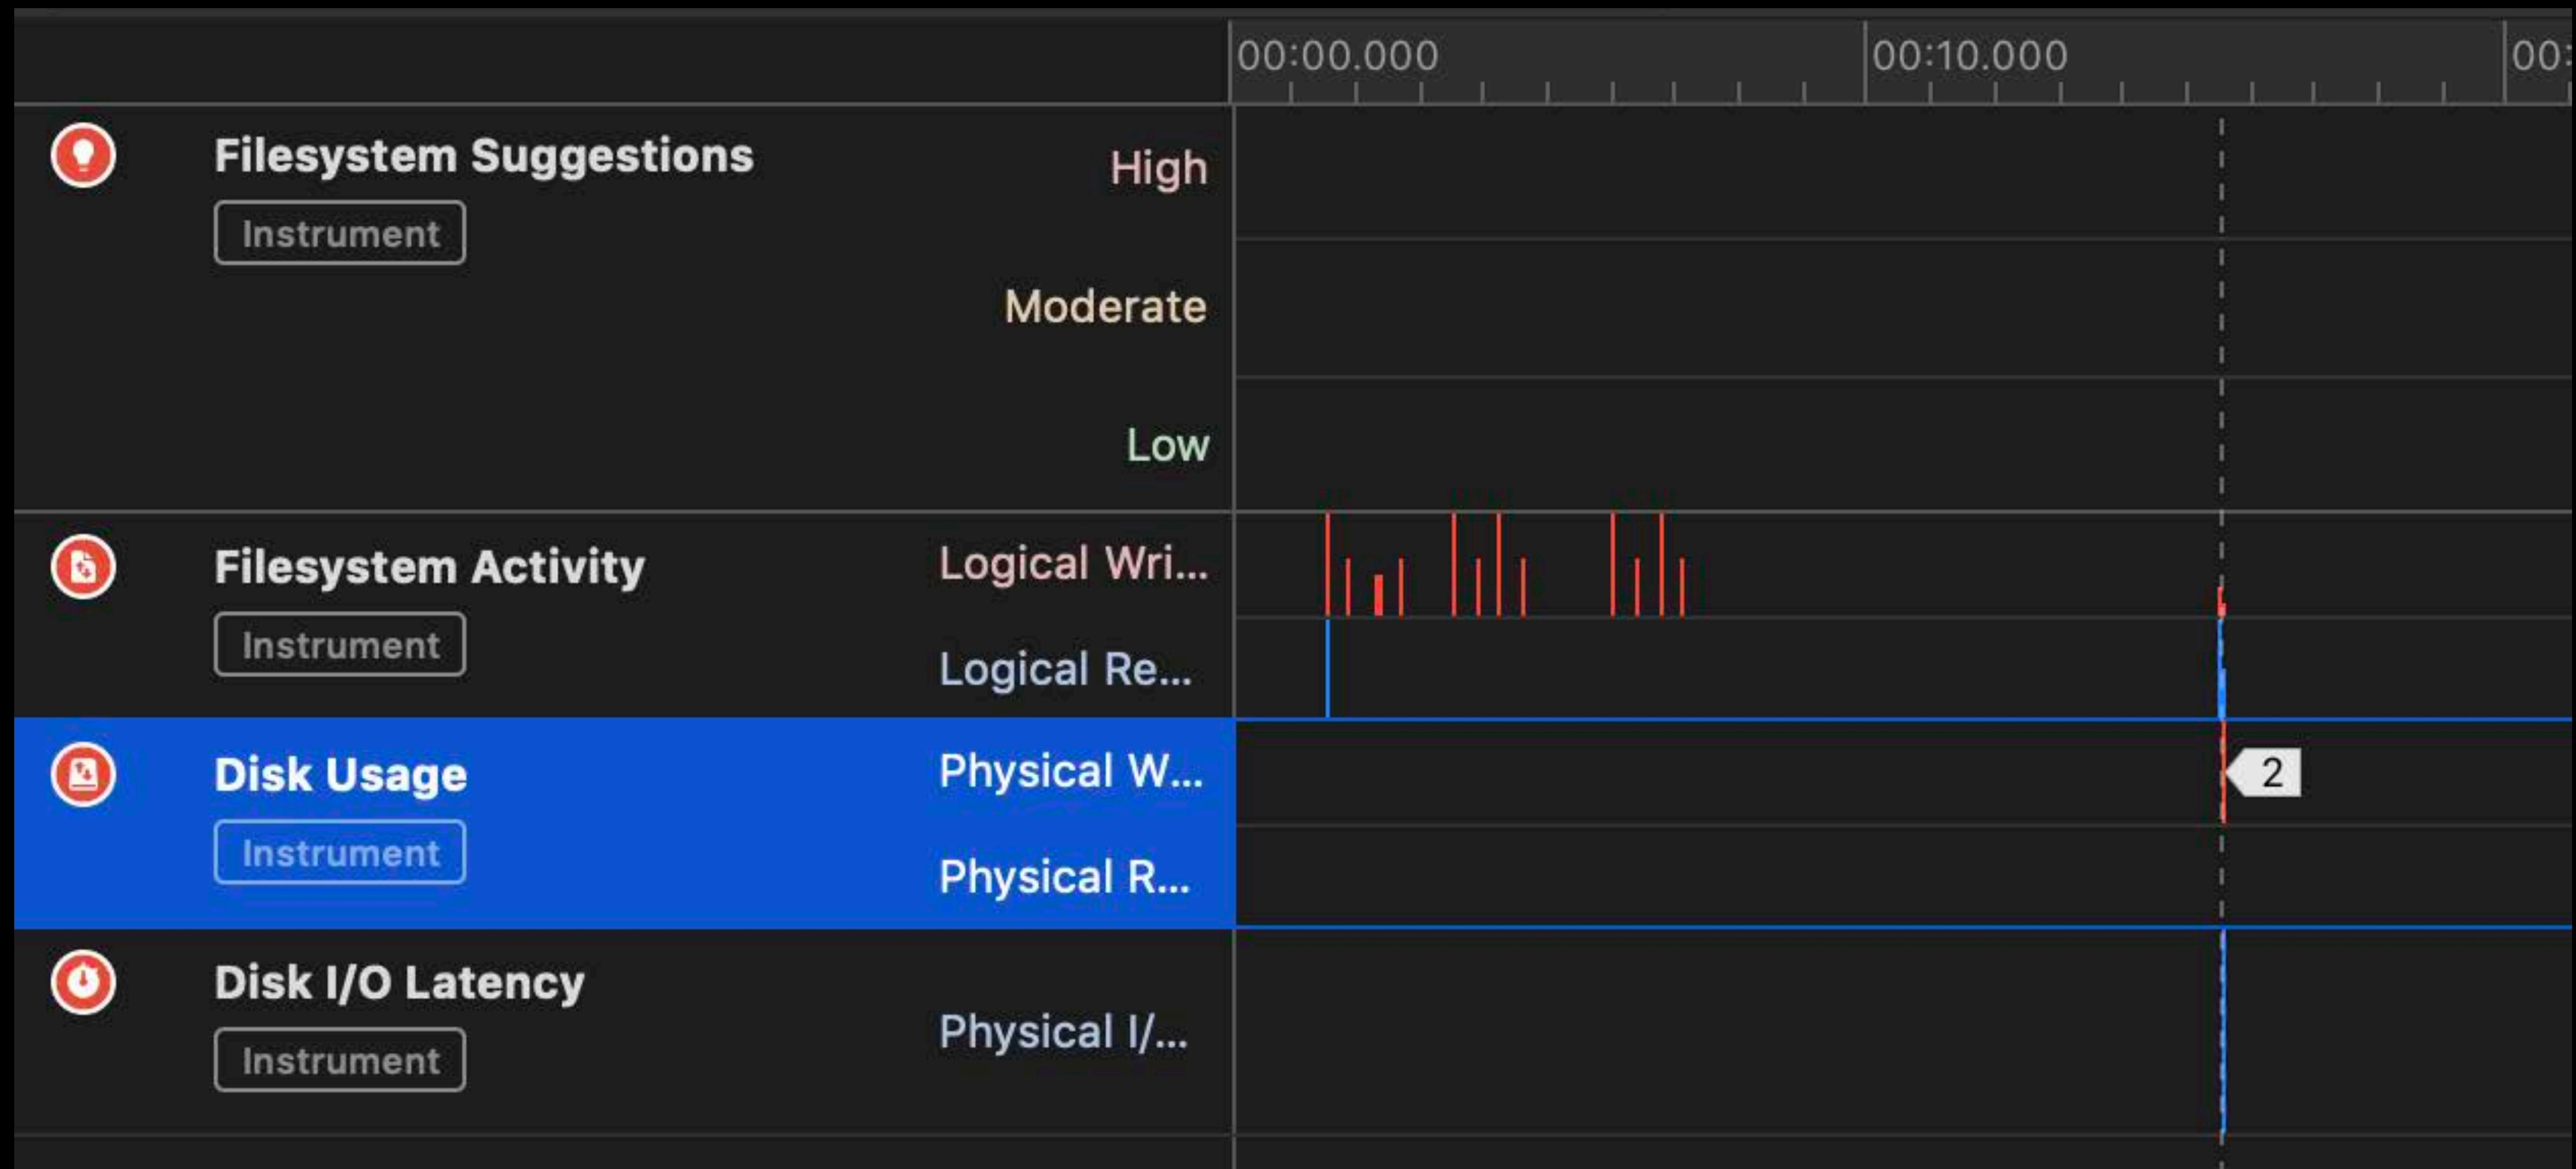
# Delete Mode Journaling

 Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	54	288.00 KiB	2.40 ms	18.98 µs	44.42 µs	57.65 µs	275.53 µs
▼ Data Write	54	288.00 KiB	2.40 ms	18.98 µs	44.42 µs	57.65 µs	275.53 µs
▼ ImageViewer (63410)	54	288.00 KiB	2.40 ms	18.98 µs	44.42 µs	57.65 µs	275.53 µs
/ImageViewer/delete_mode.db-journal	36	192.00 KiB	1.96 ms	18.98 µs	54.33 µs	68.71 µs	275.53 µs
/ImageViewer/delete_mode.db	18	96.00 KiB	443.16 µs	19.35 µs	24.62 µs	4.24 µs	35.66 µs

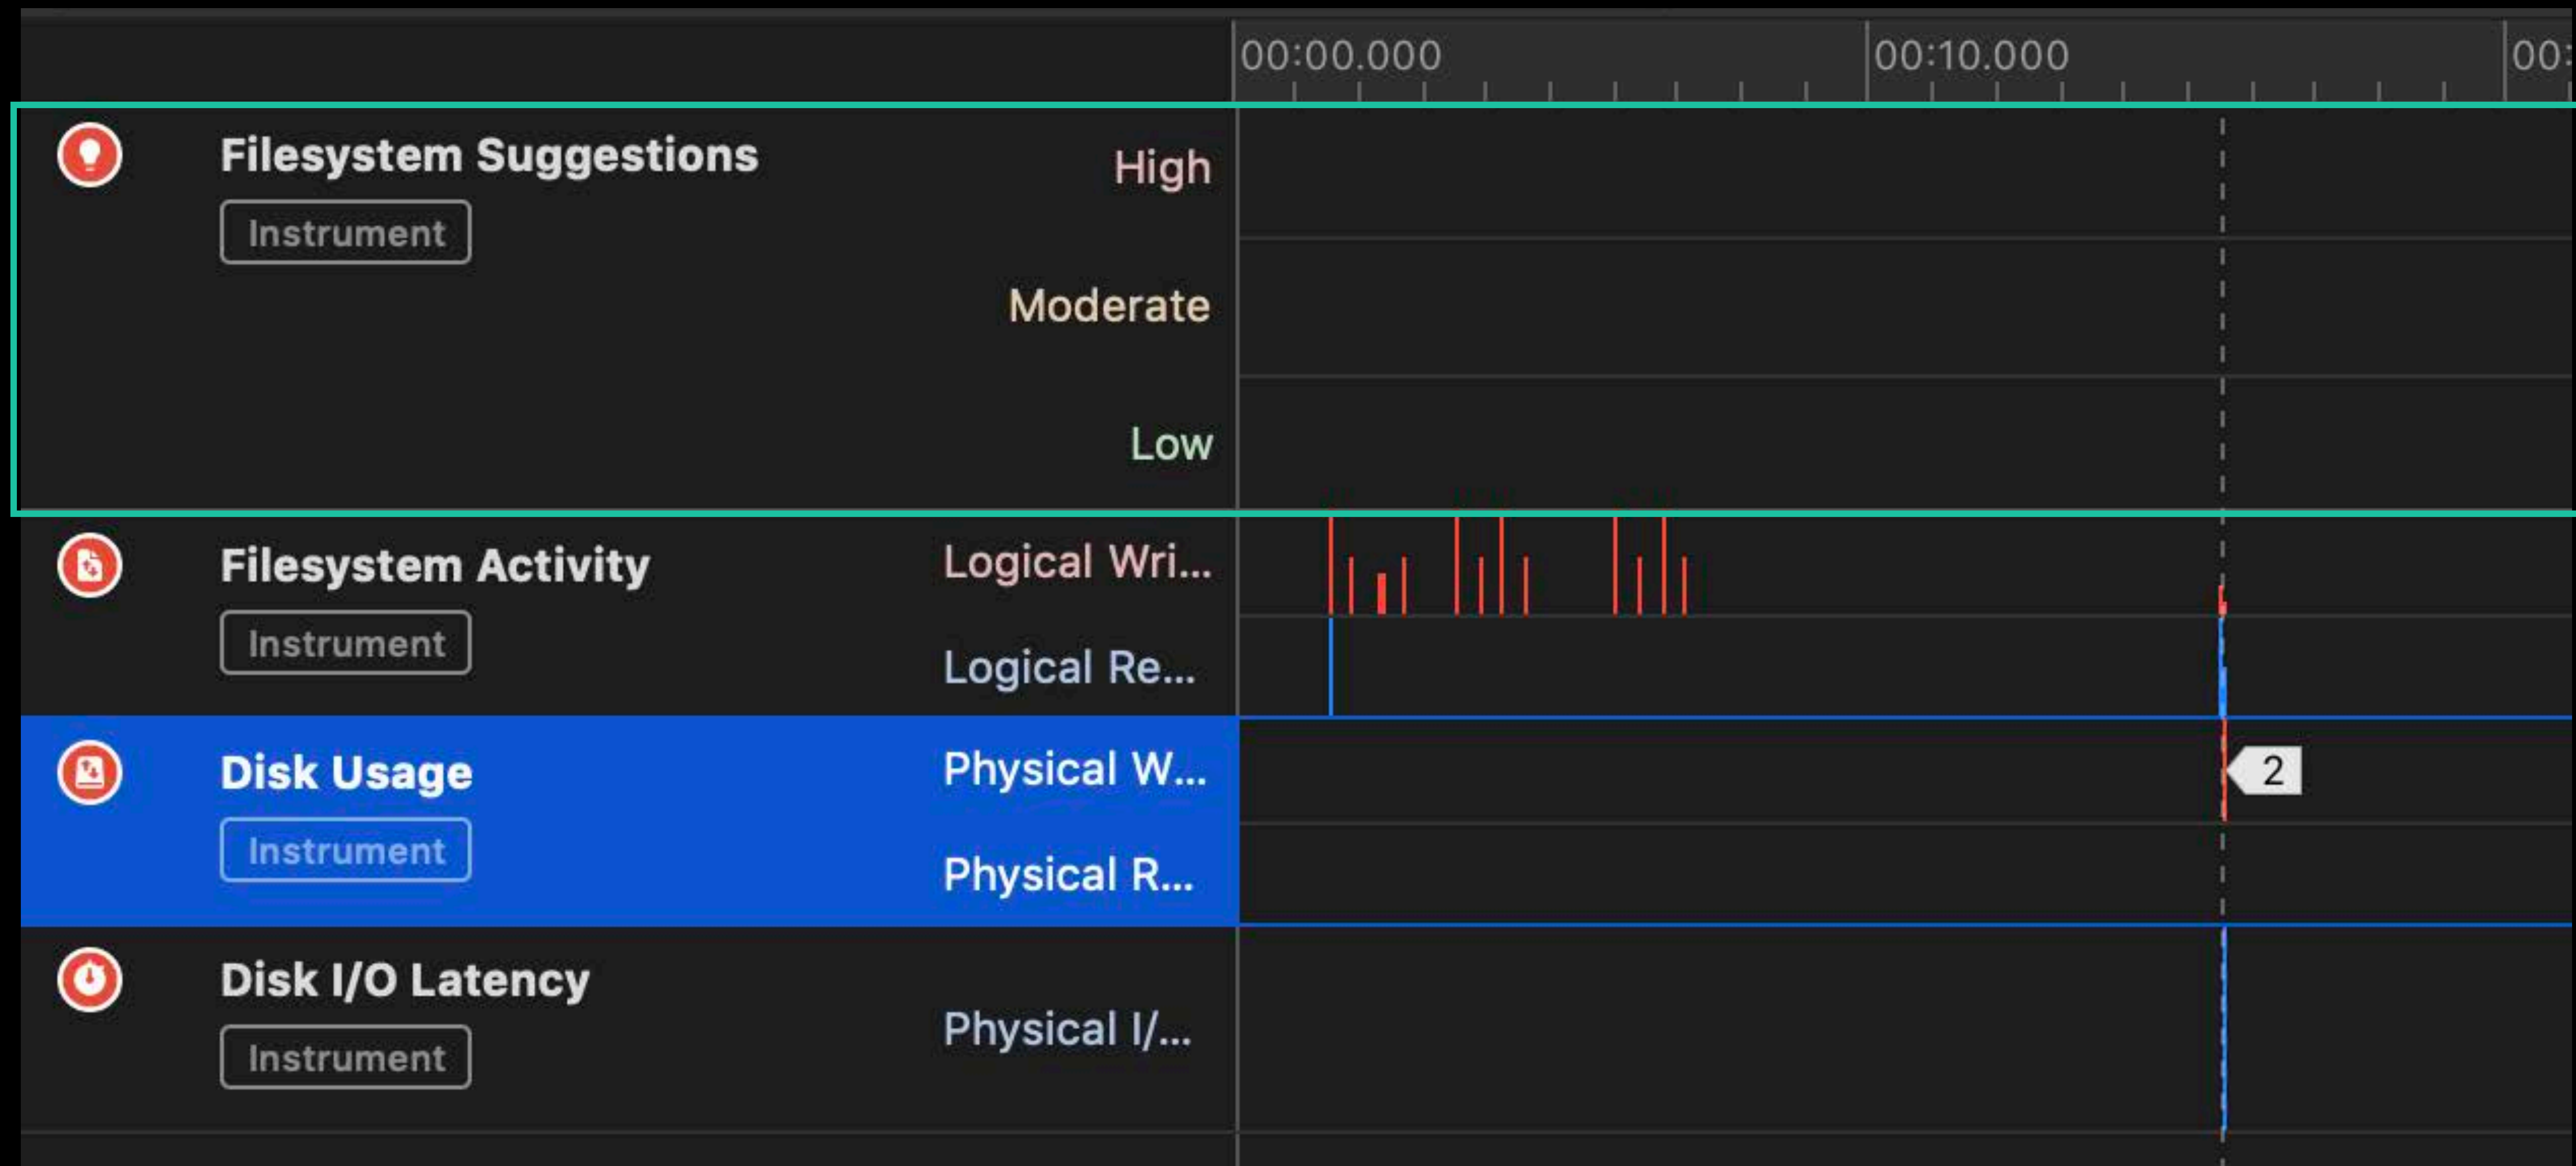


# WAL Mode Journaling



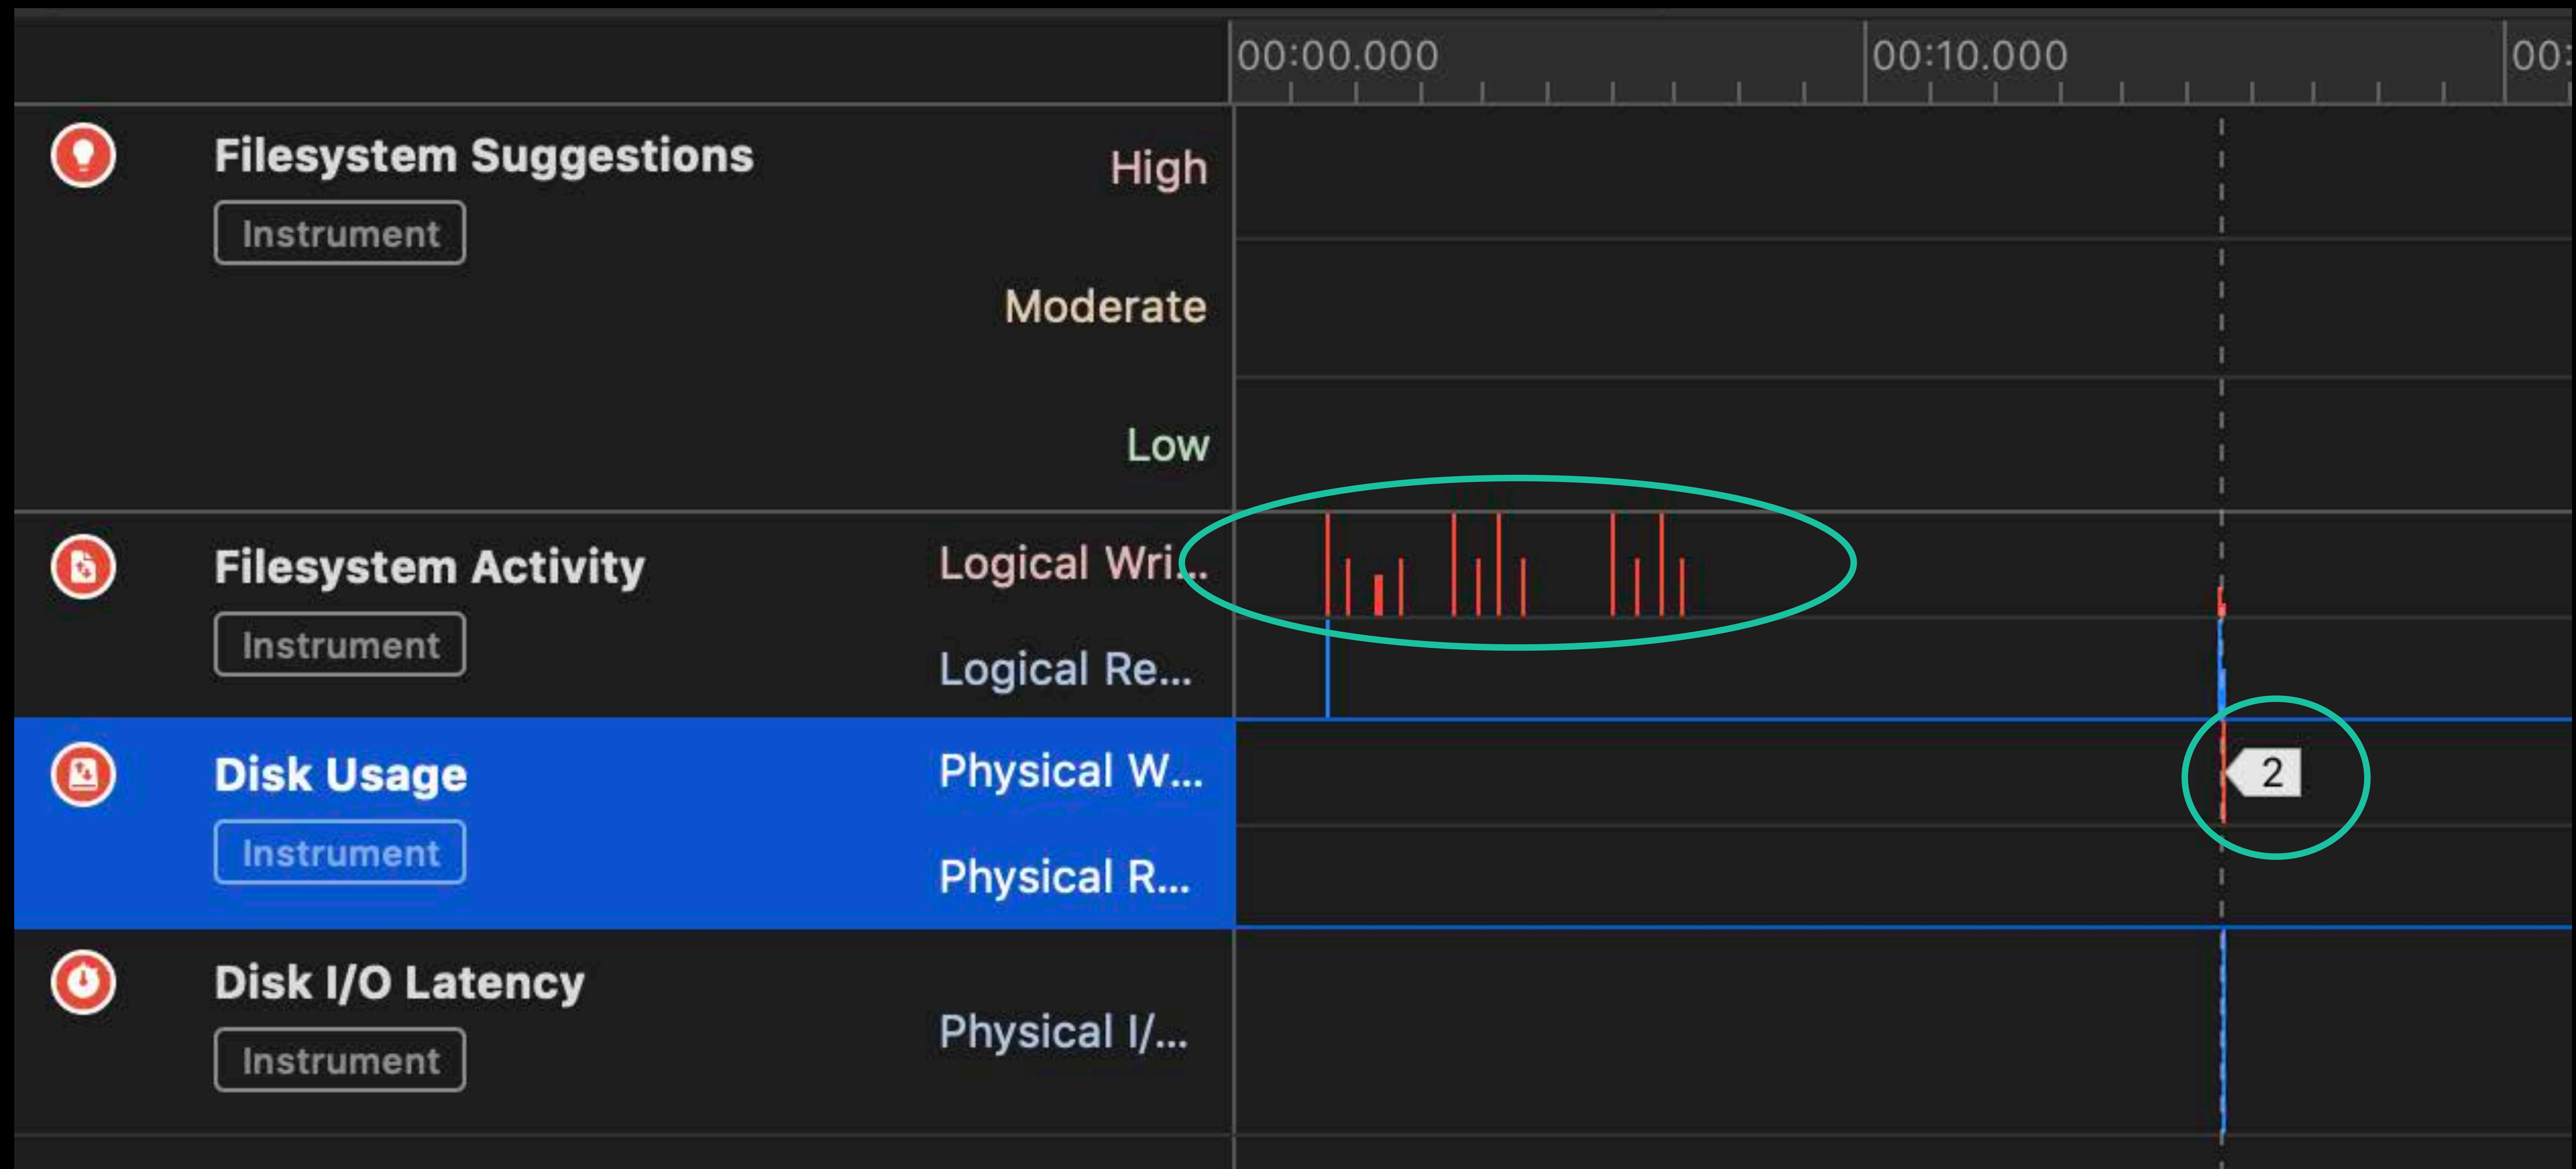


# WAL Mode Journaling





# WAL Mode Journaling

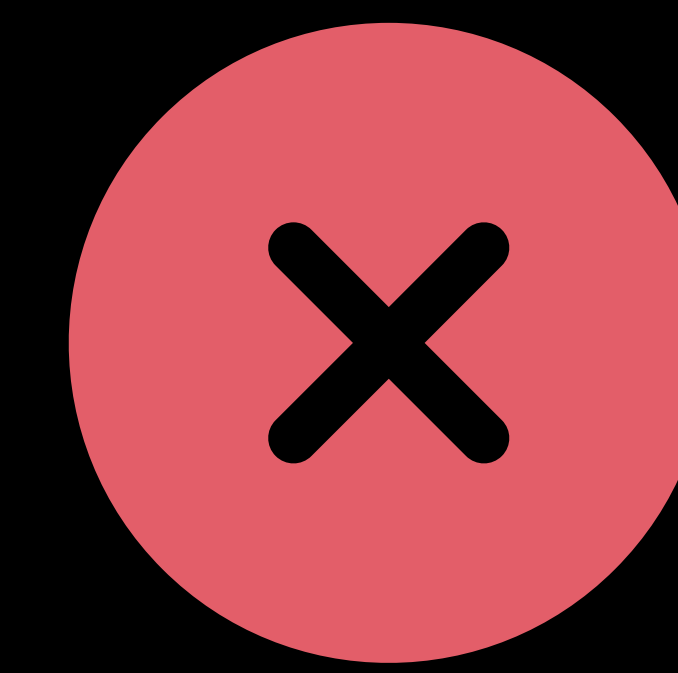
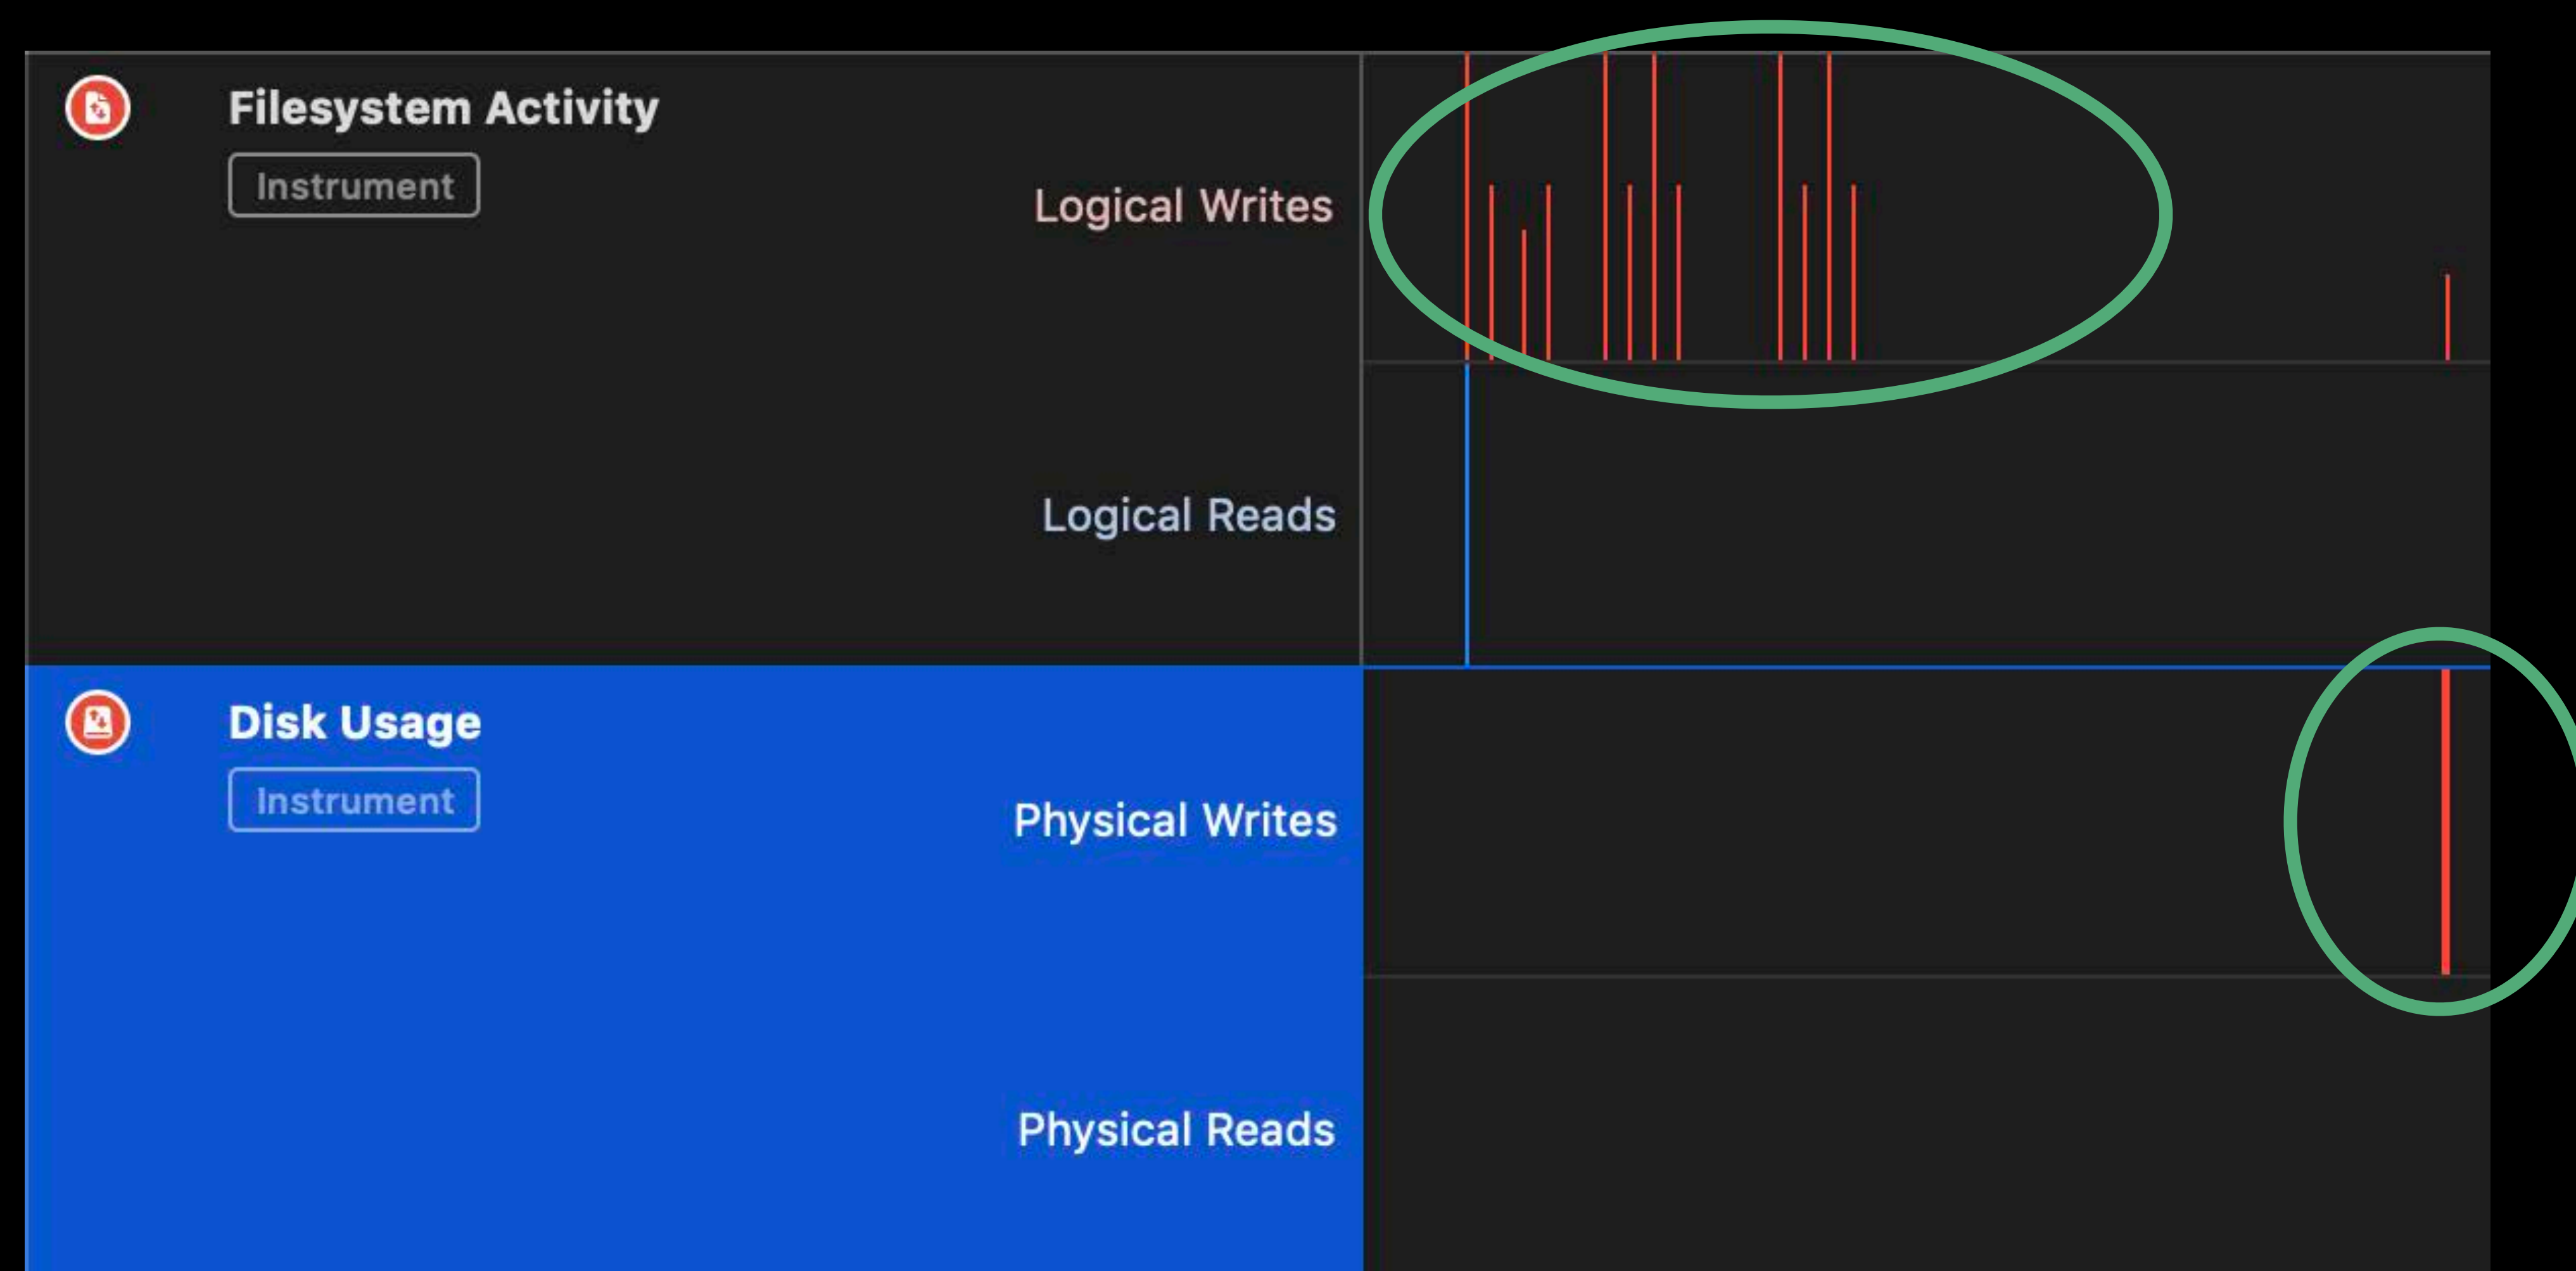




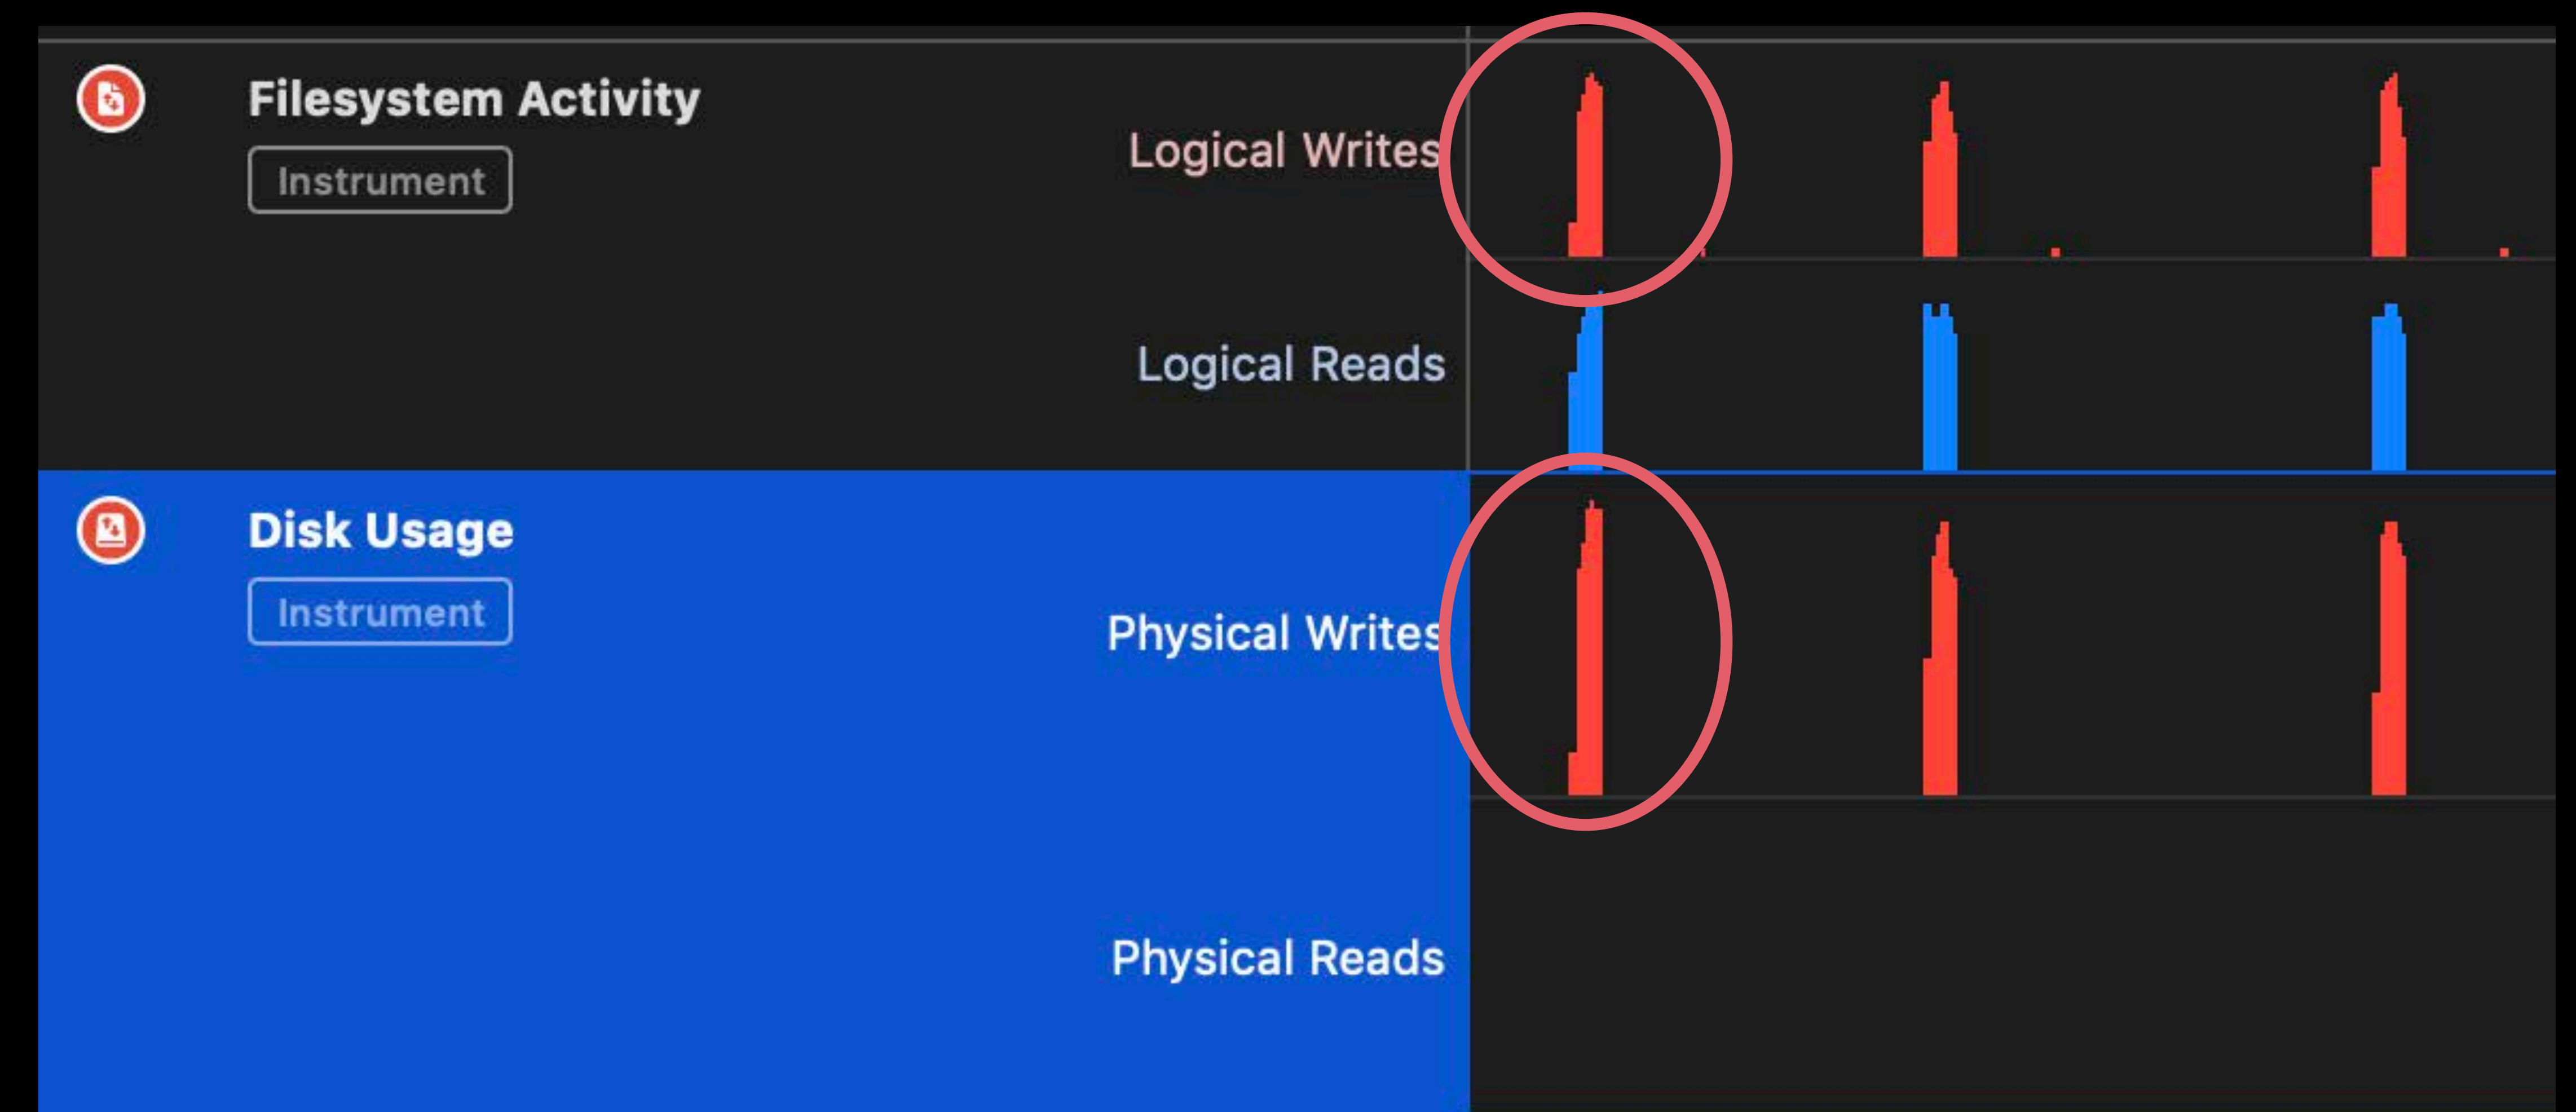
# Journaling Comparison



WAL Journaling



Delete Mode Journaling





# Journaling Mode Comparison



WAL Journaling



Delete Mode Journaling

Filesystem Activity > Filesystem Statistics	
Syscall / Process	Count
▼ * All *	151
▶ fsync	1
▶ ftruncate	2
▶ stat64	2
▶ pread	5
▶ fstat64	5
▶ pwrite	9
▶ write_nocancel	62
▶ fcntl	65

Filesystem Activity > Filesystem Statistics	
Syscall / Process	Count
▼ * All *	240
▶ fstatfs64	4
▶ unlink	4
▶ pread	12
▶ fstat64	16
▶ fsync	16
▶ stat64	24
▶ pwrite	40
▶ write_nocancel	40
▶ fcntl	84

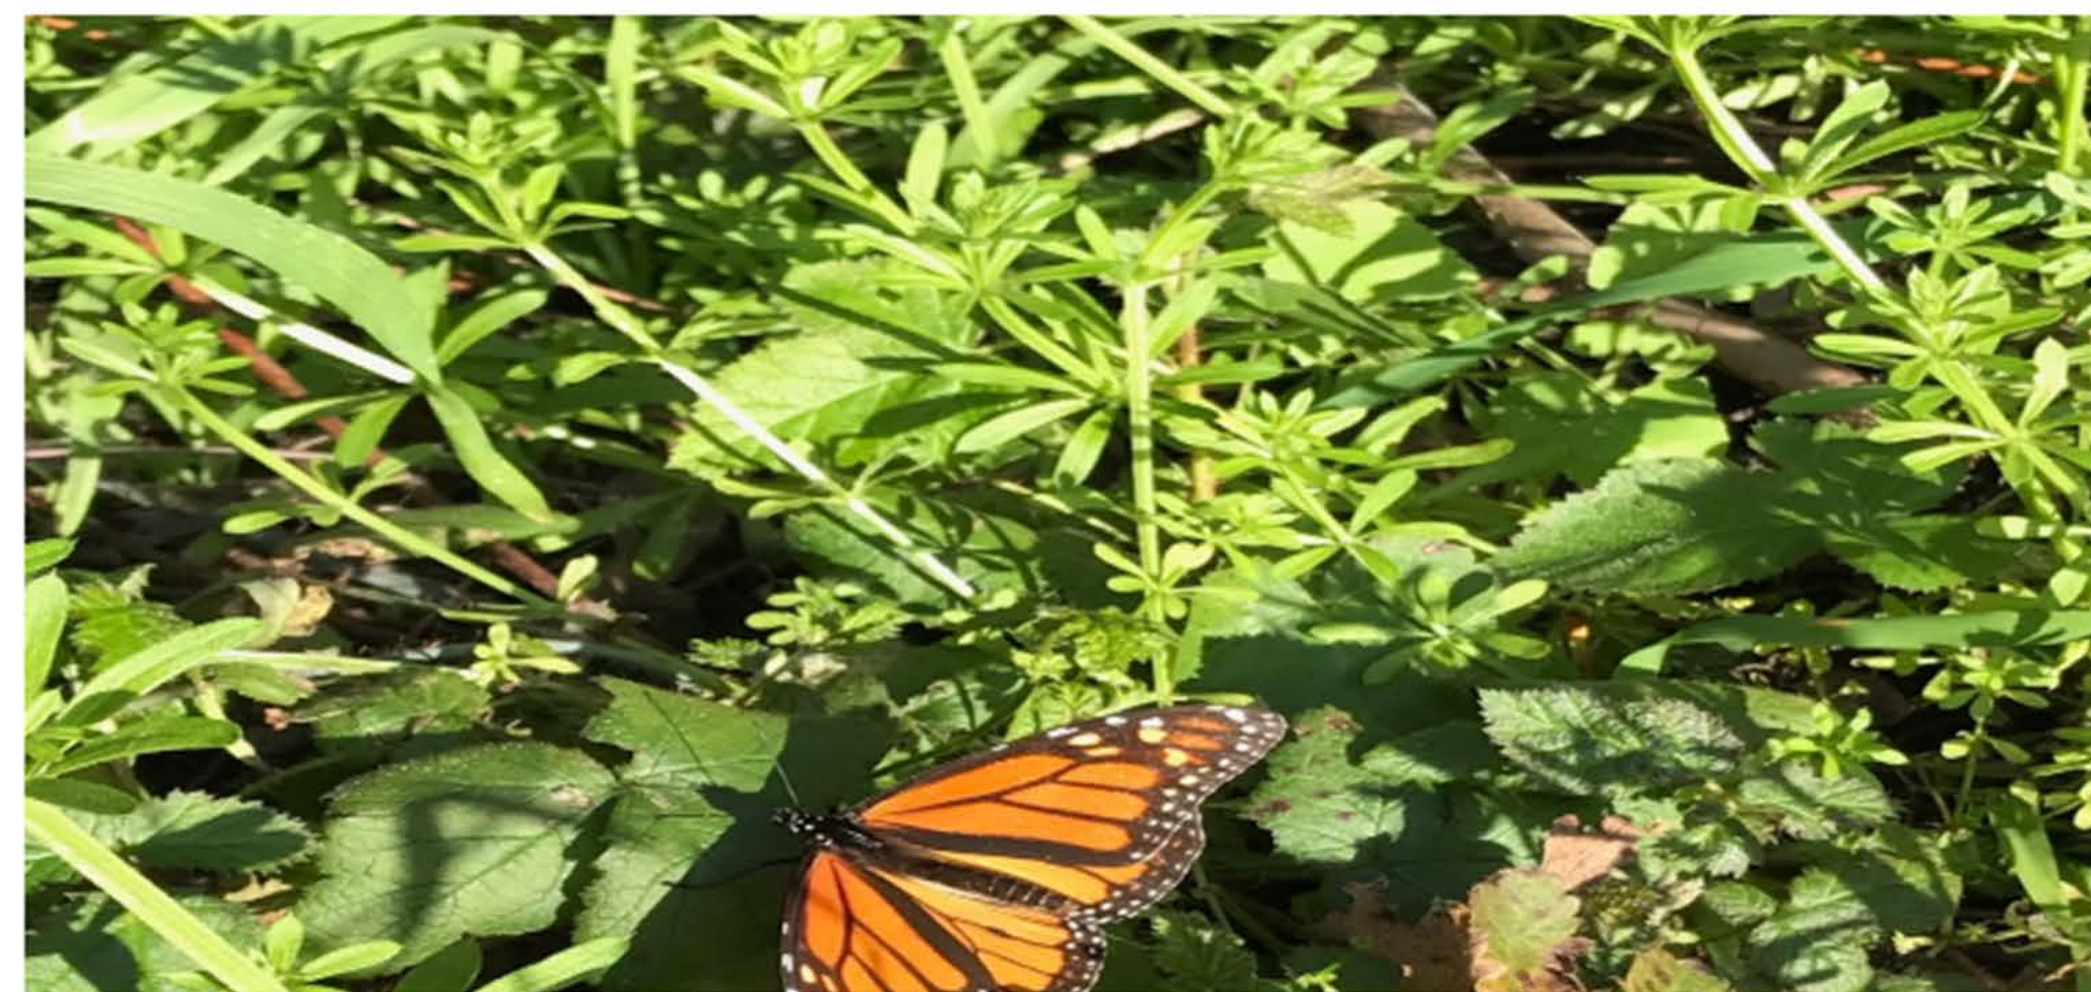
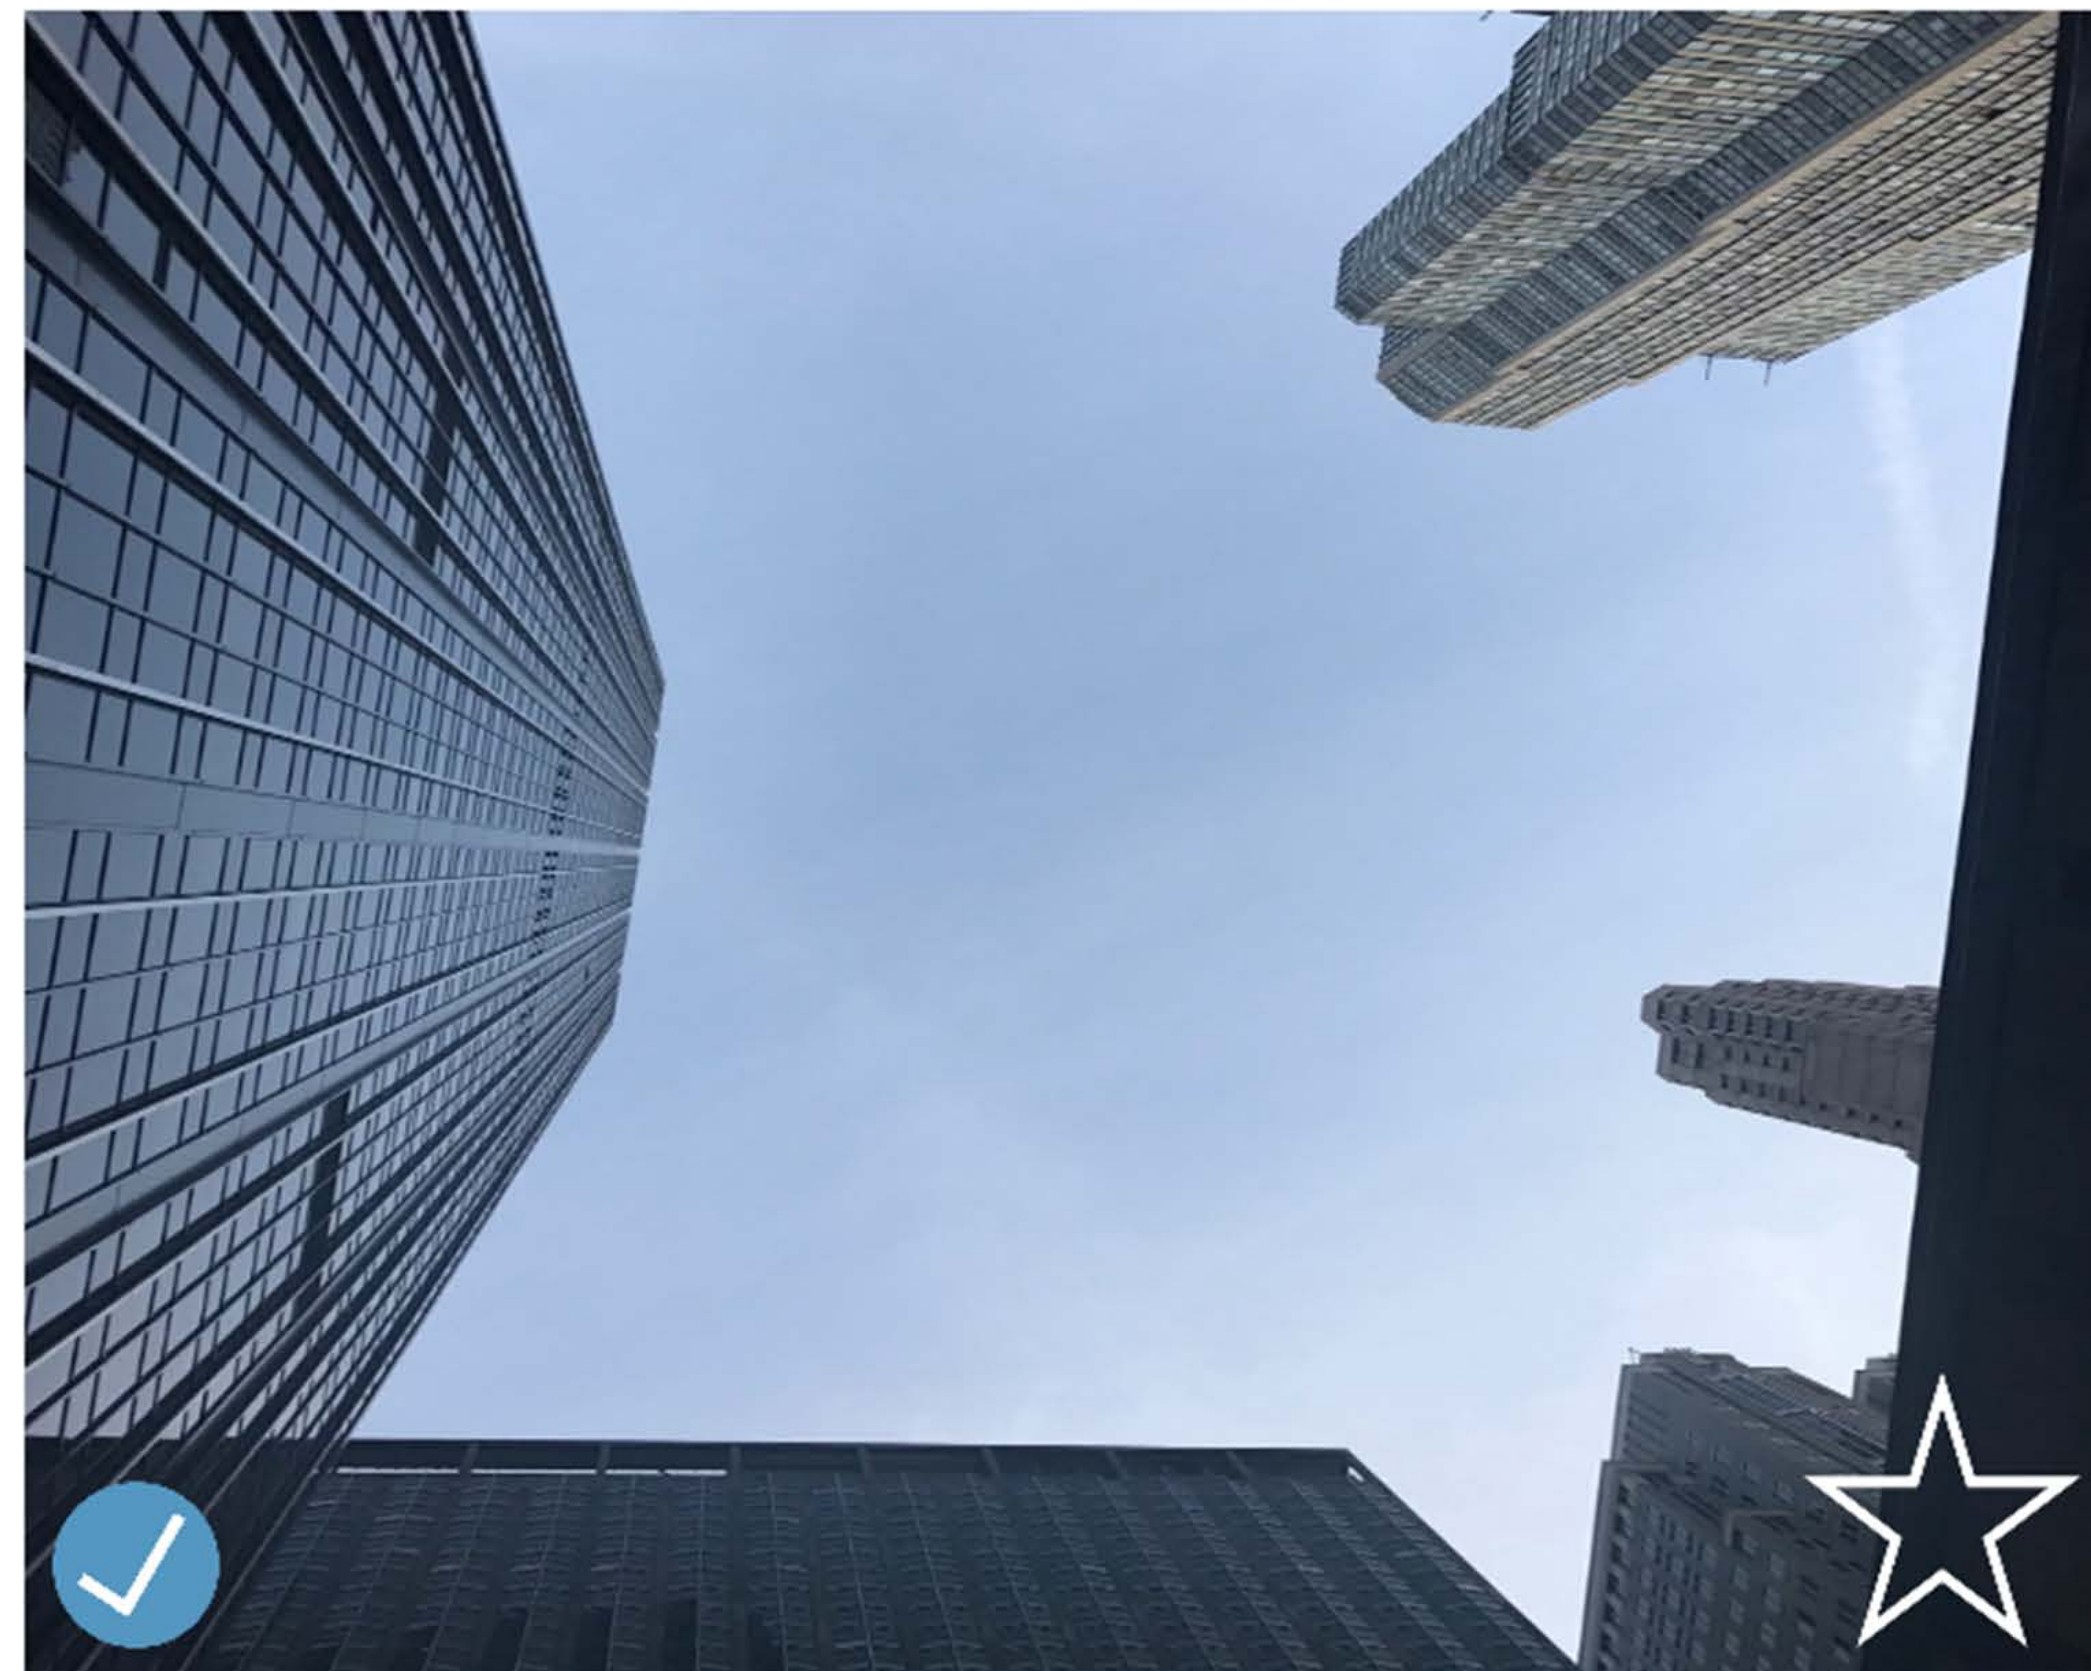


9:41



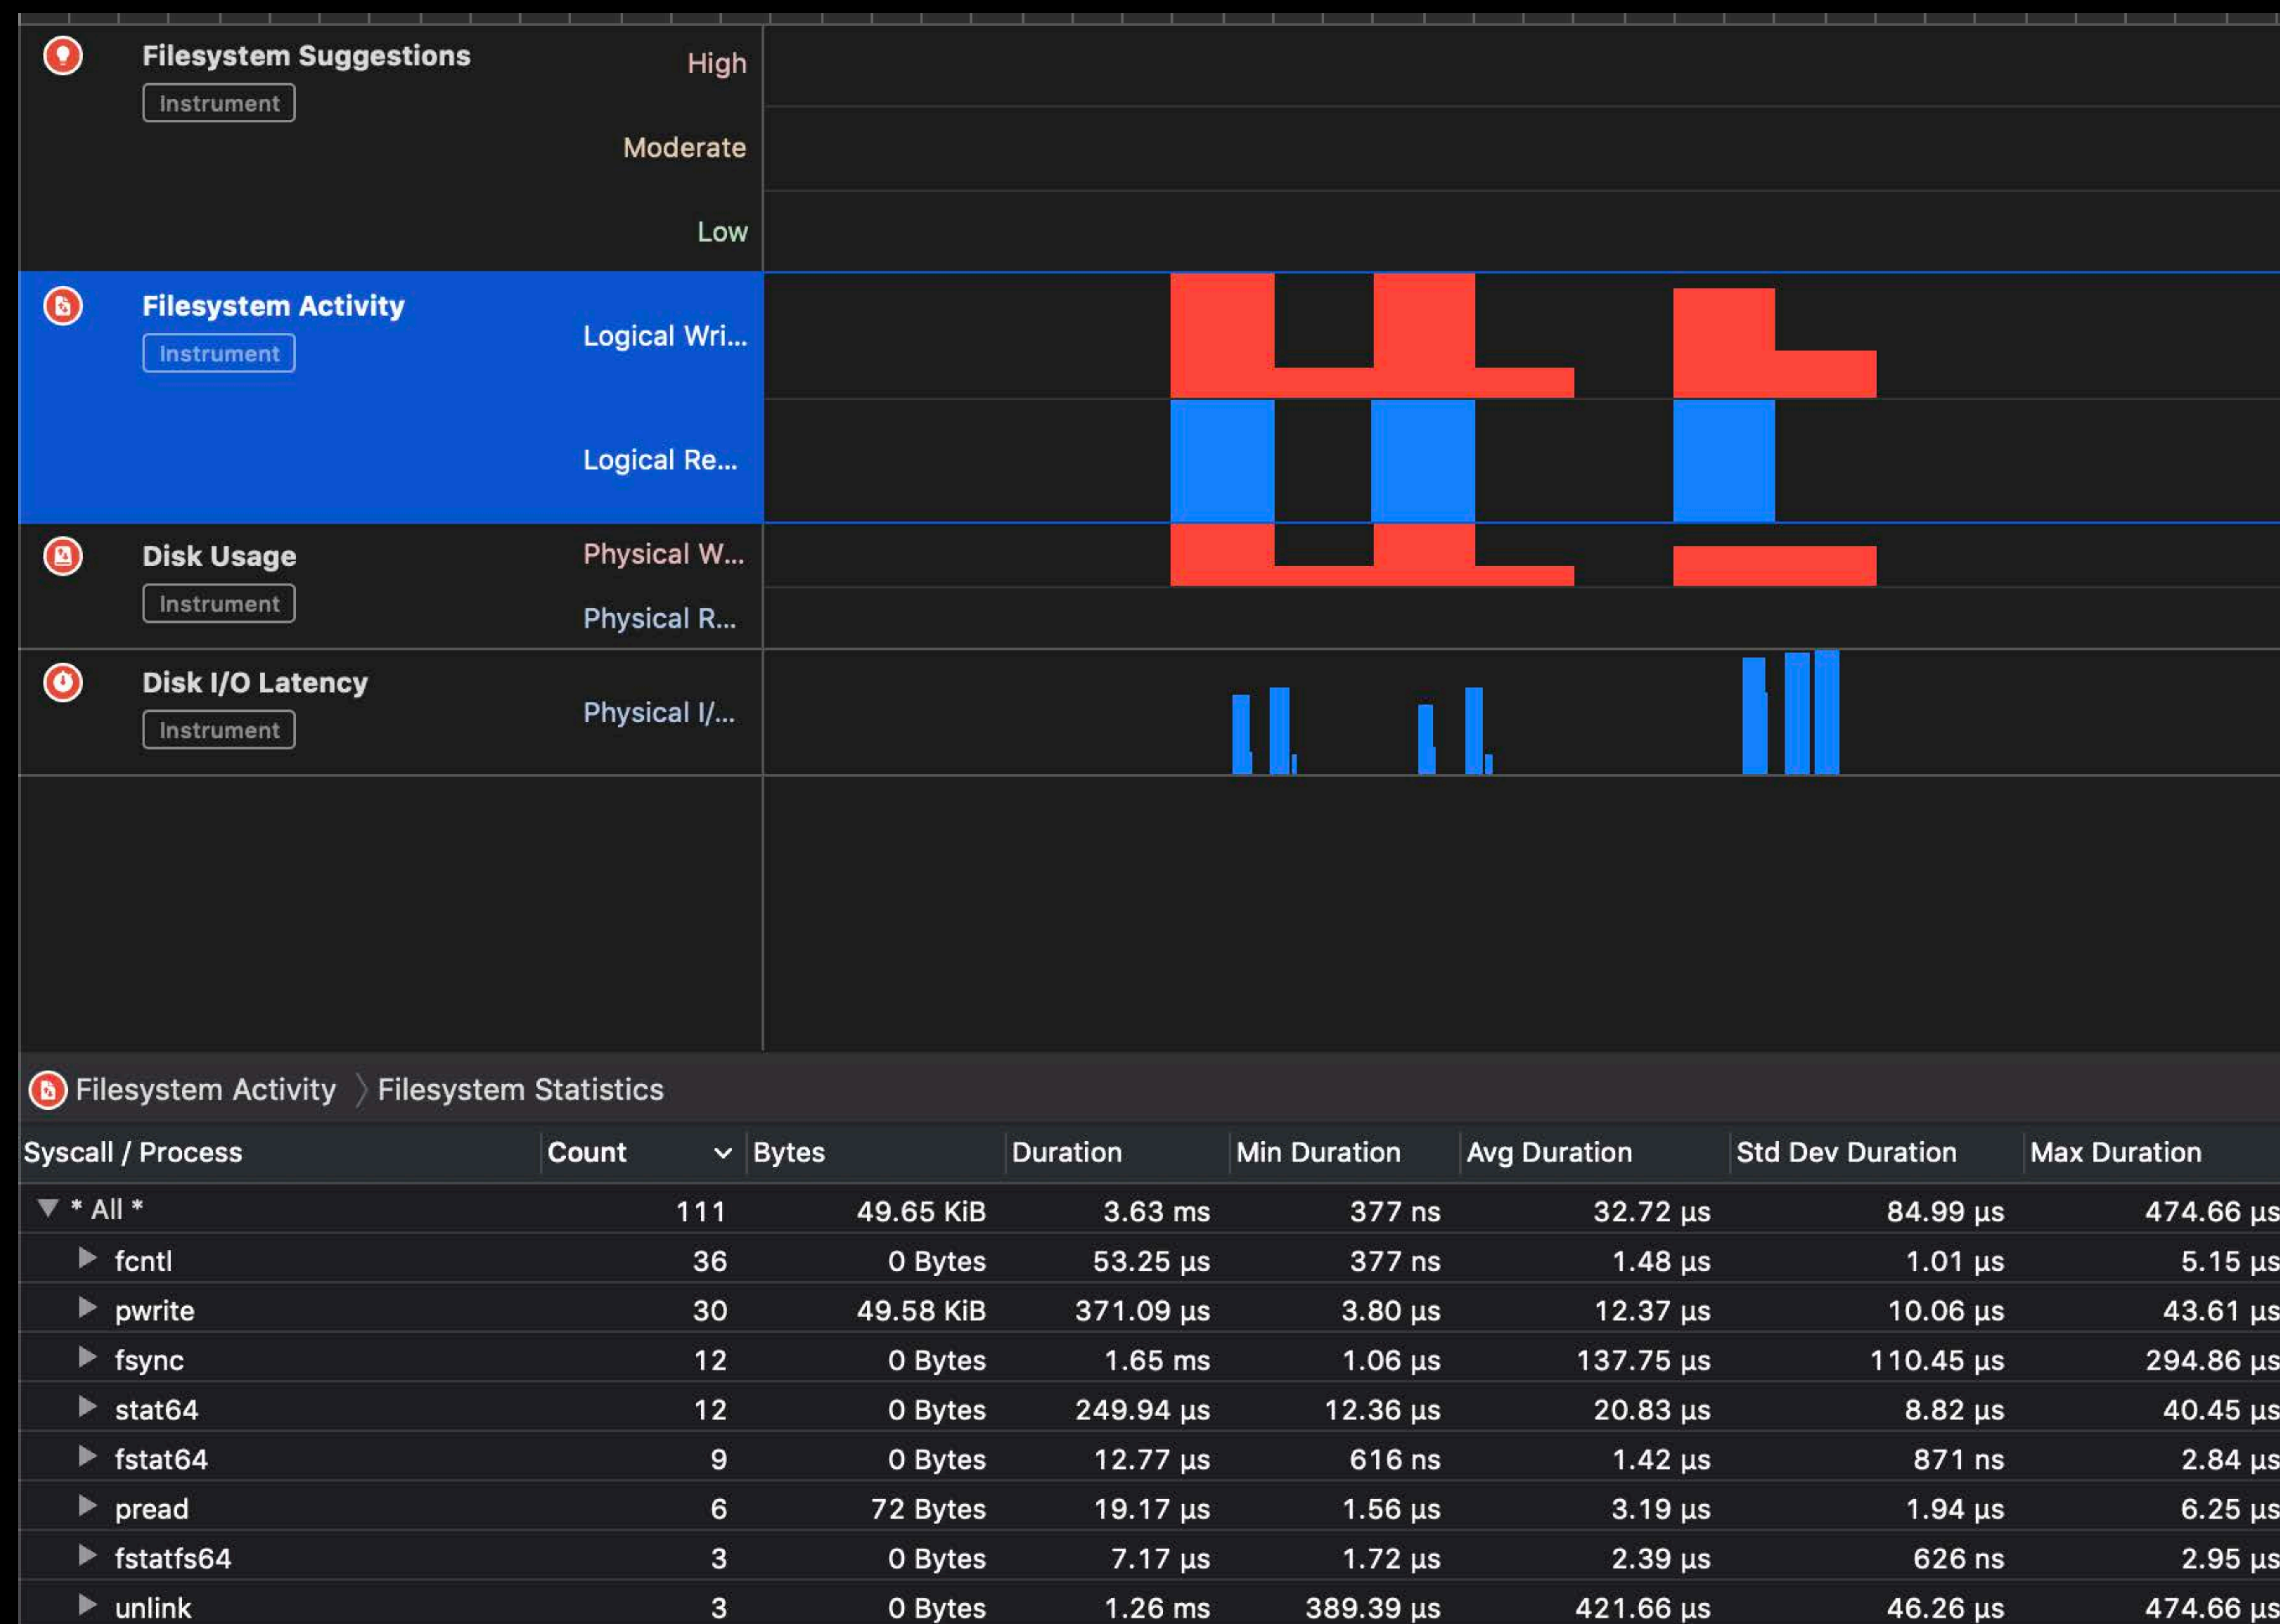
Add

Delete



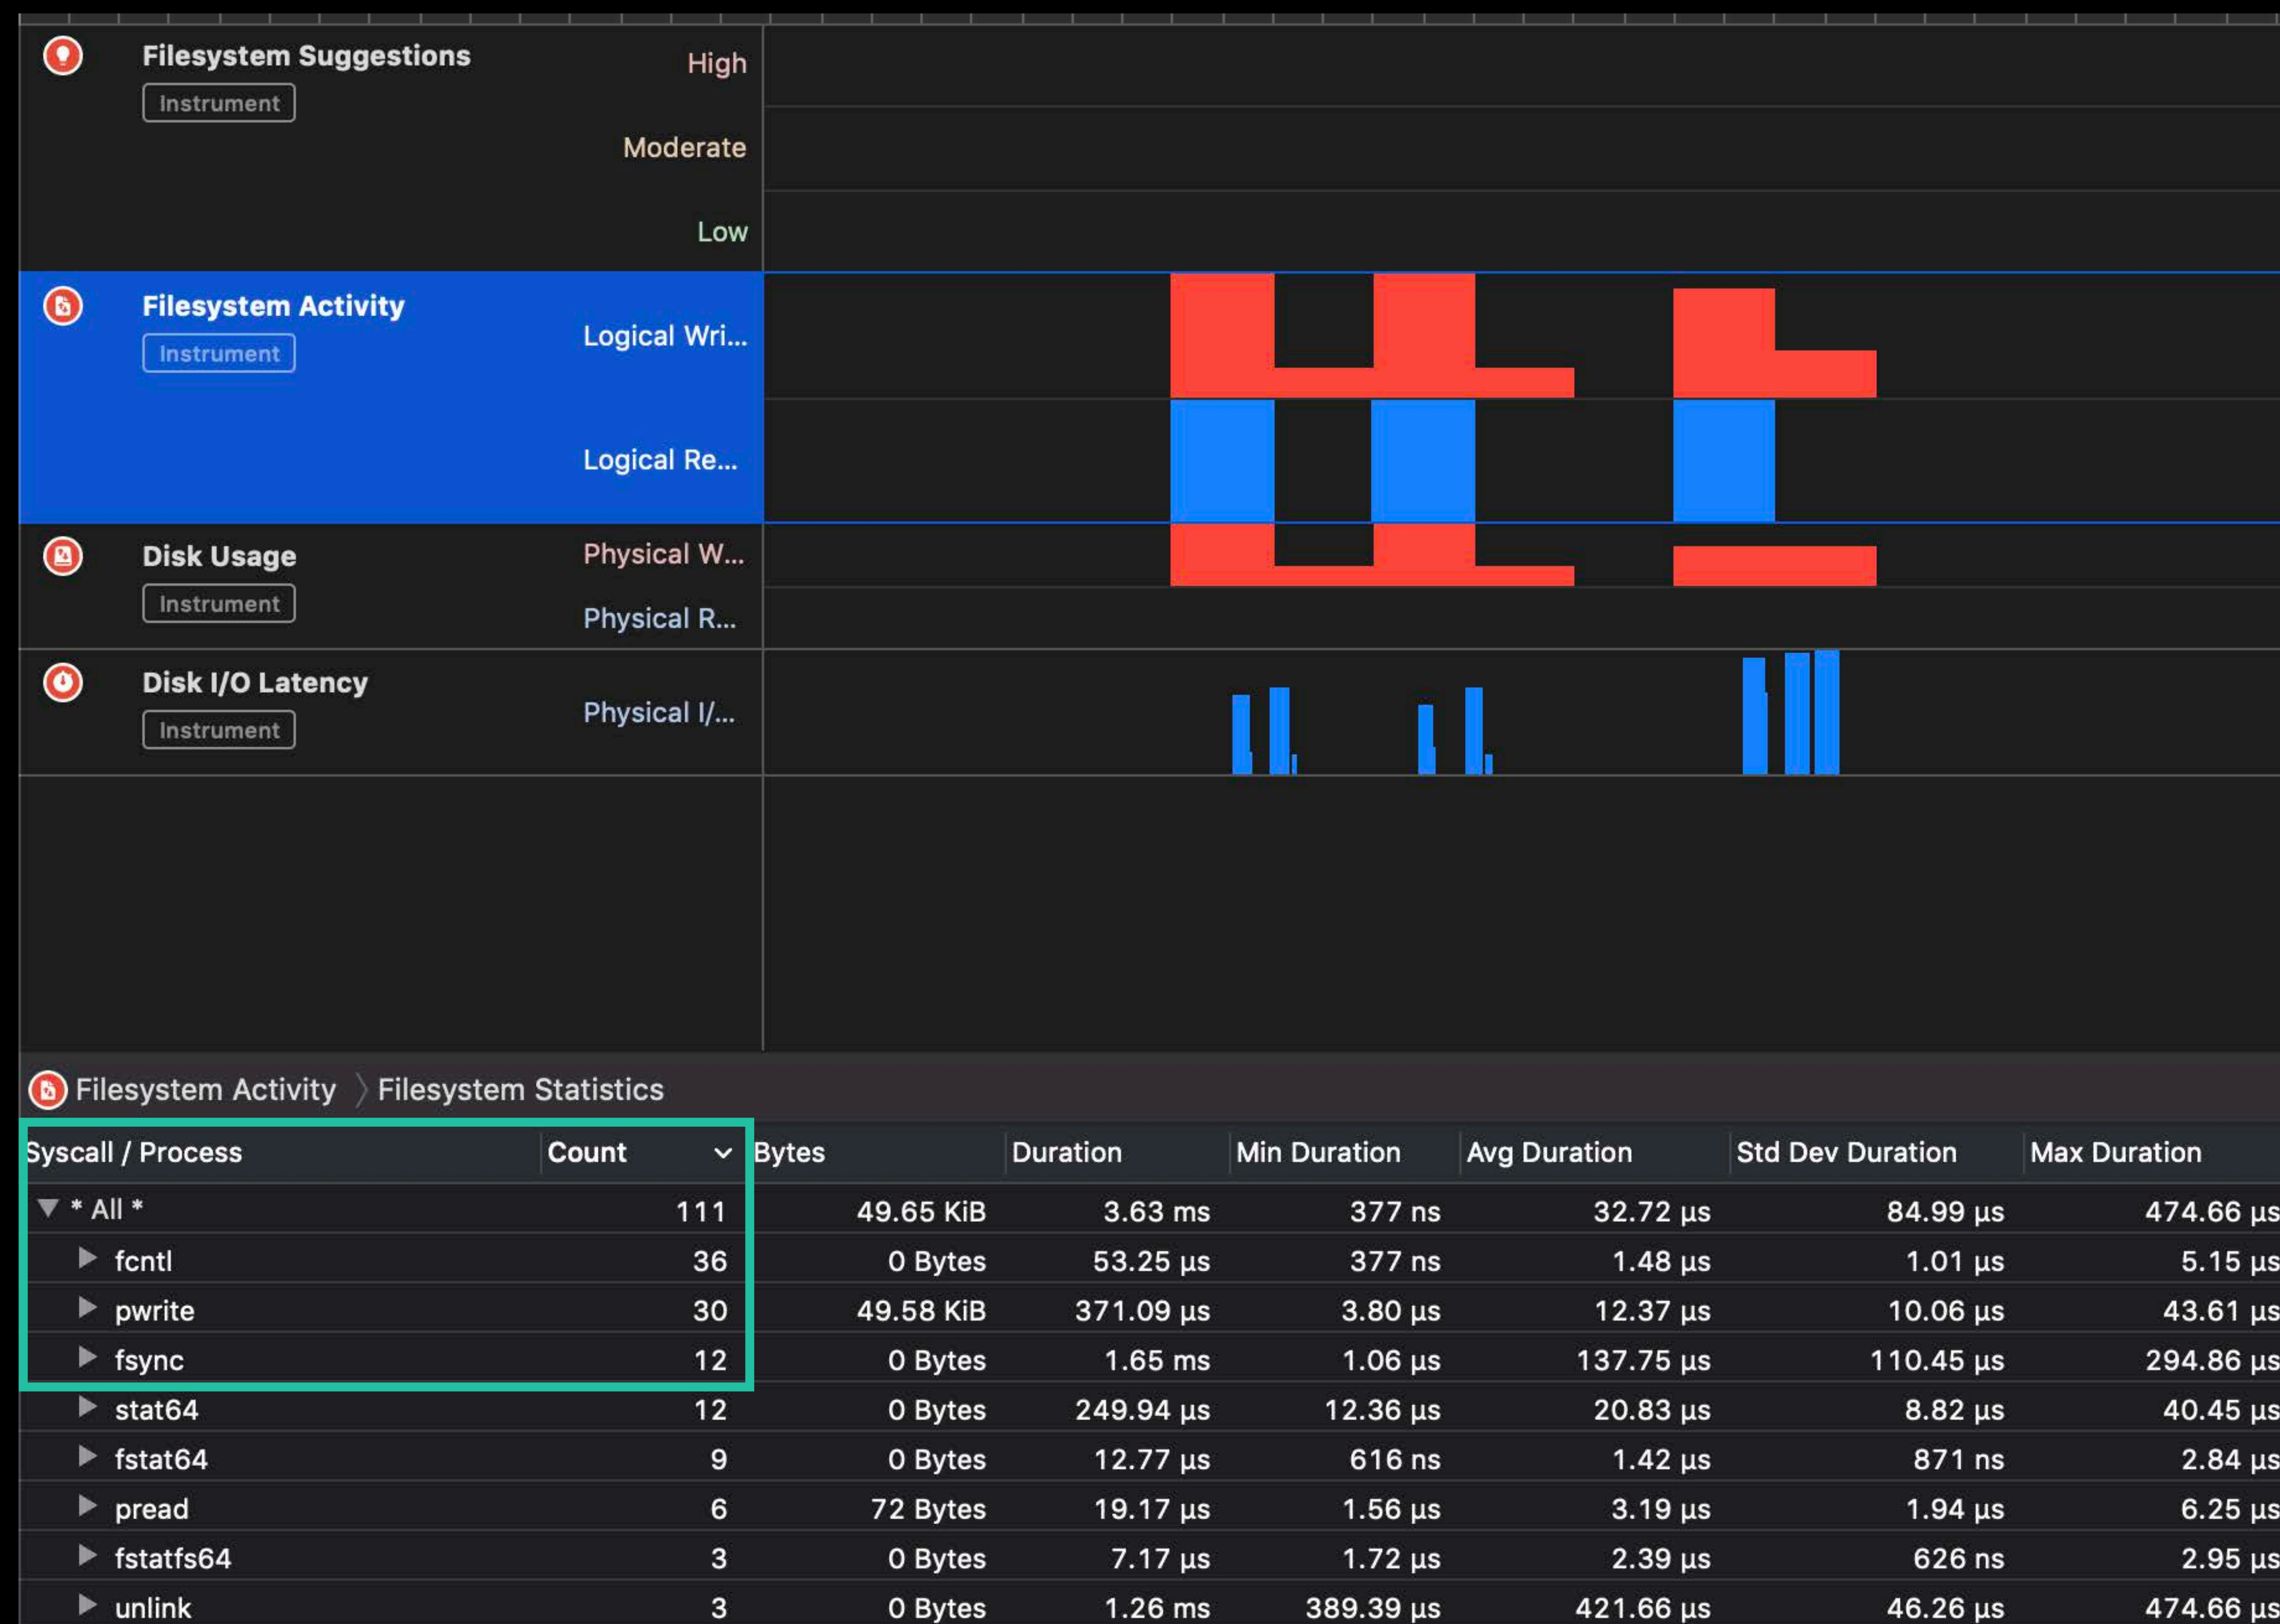


# Single Statement Transactions



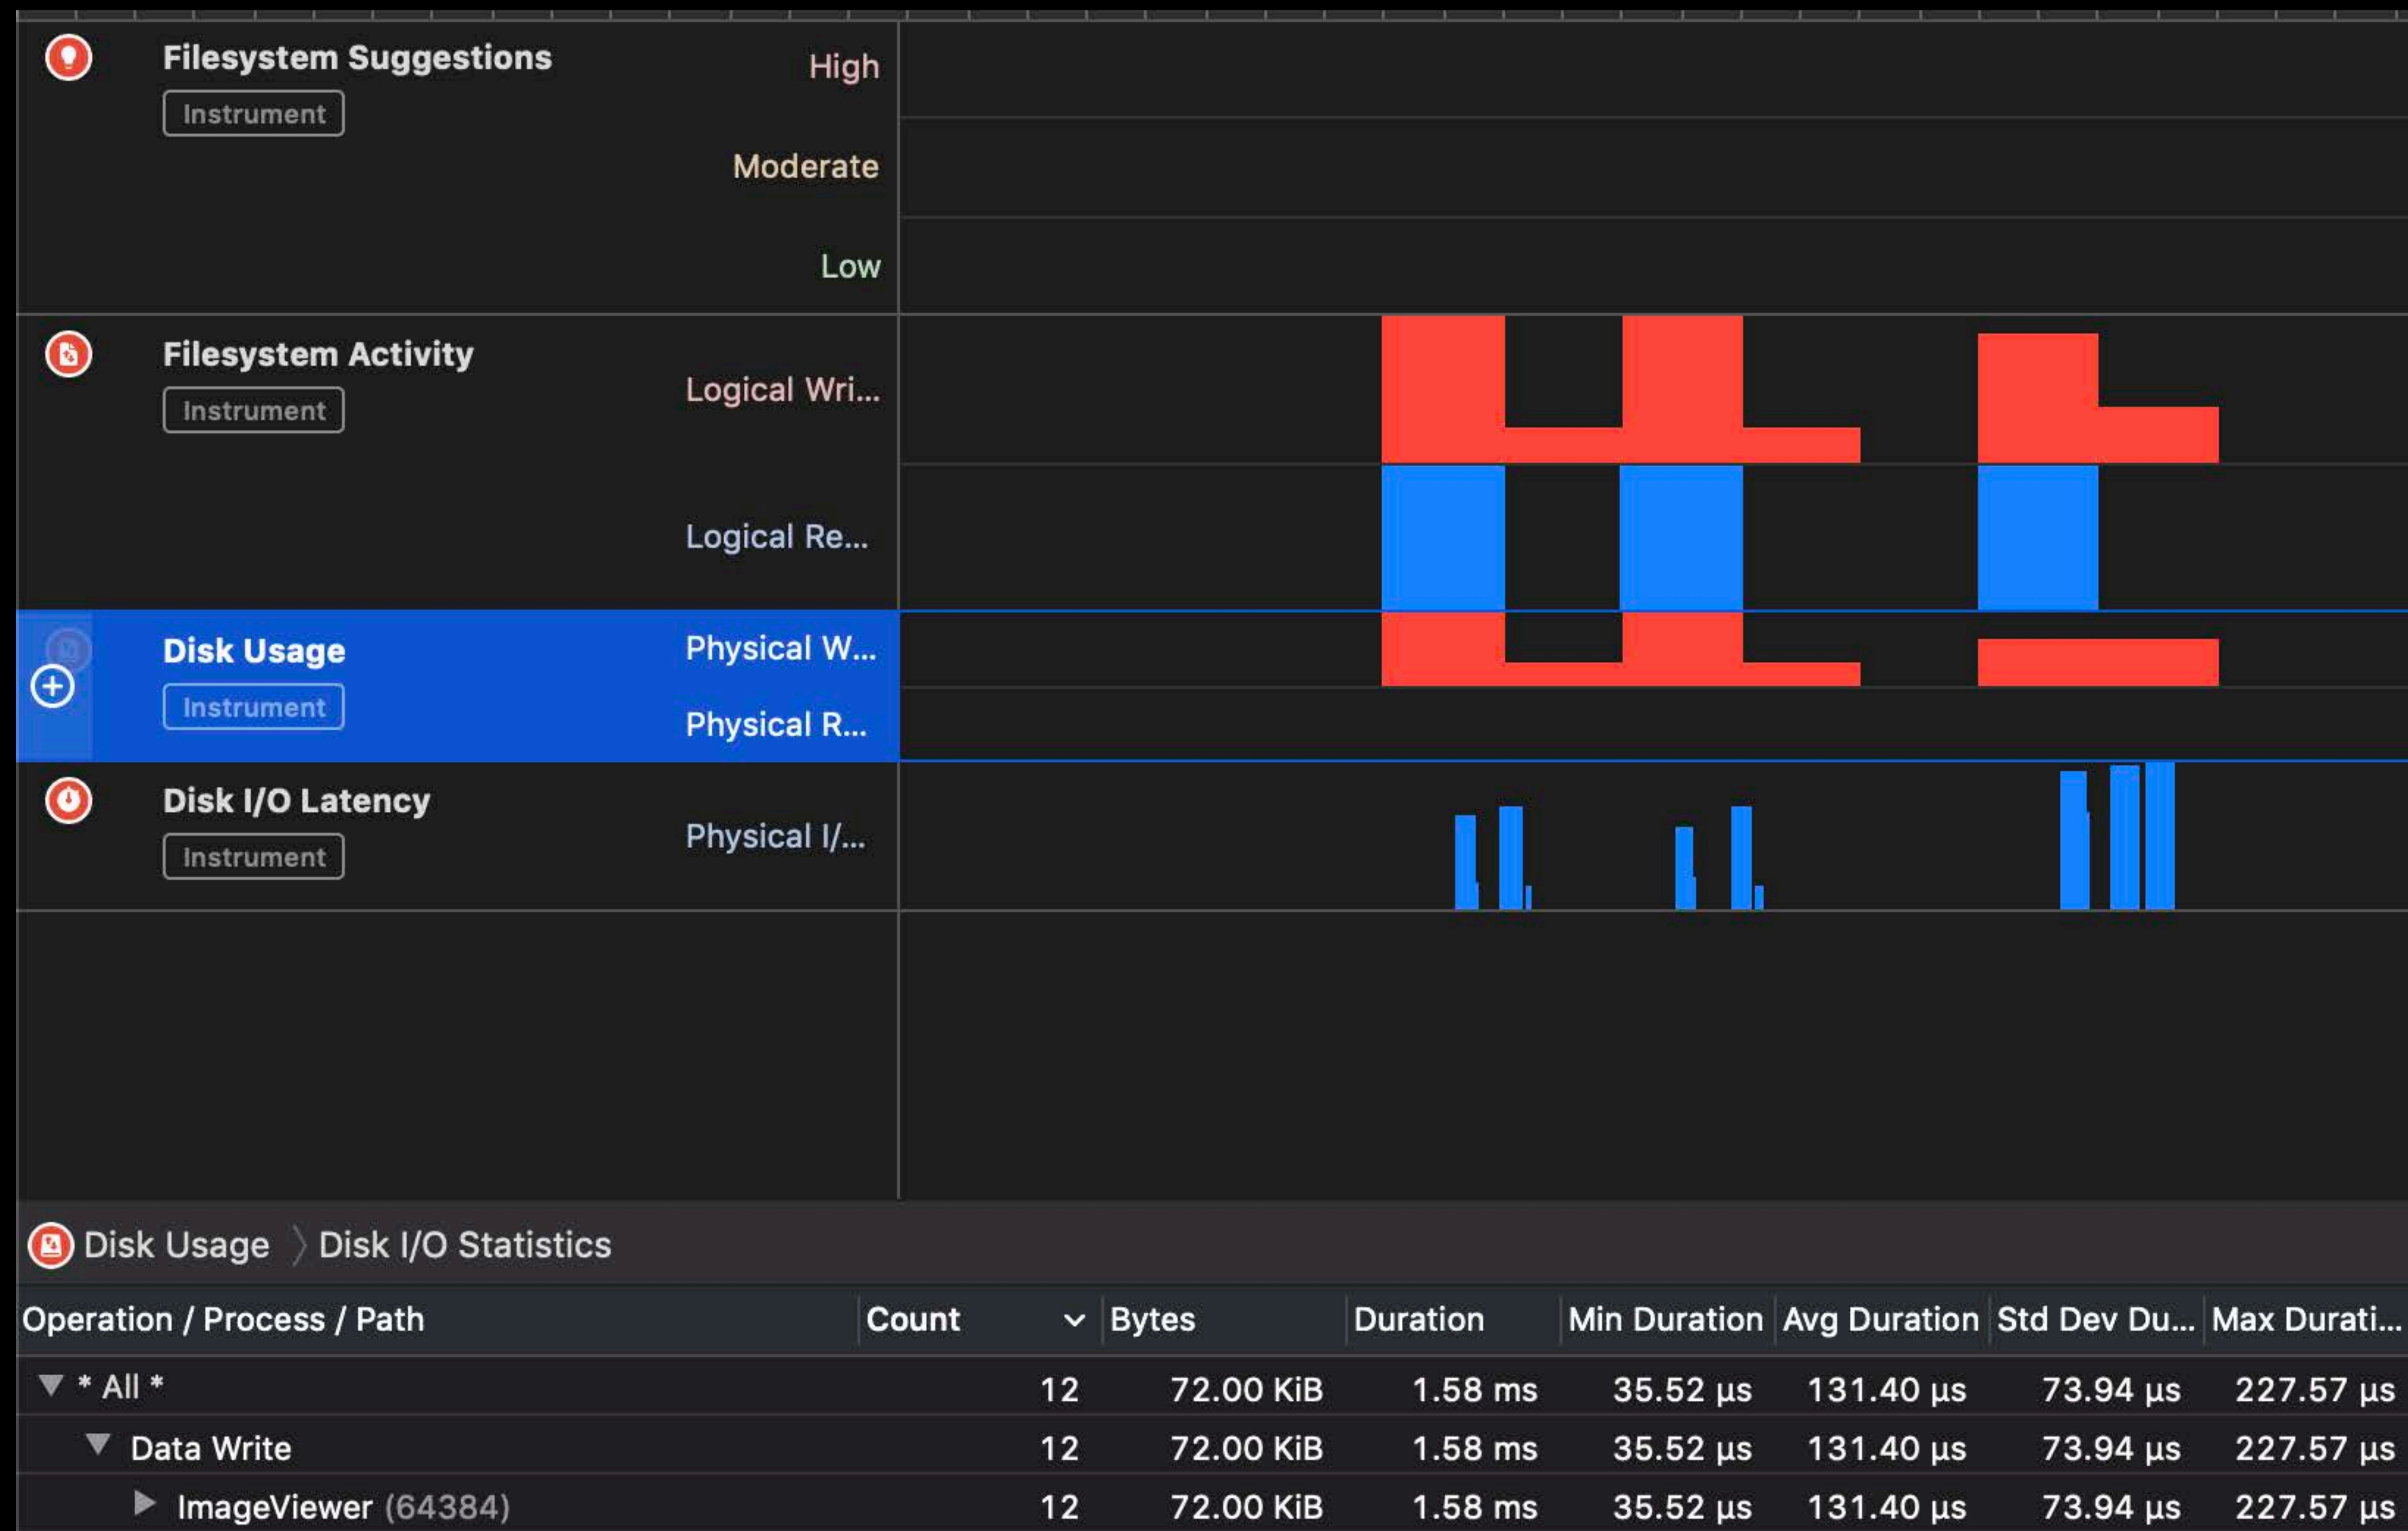


# Single Statement Transactions



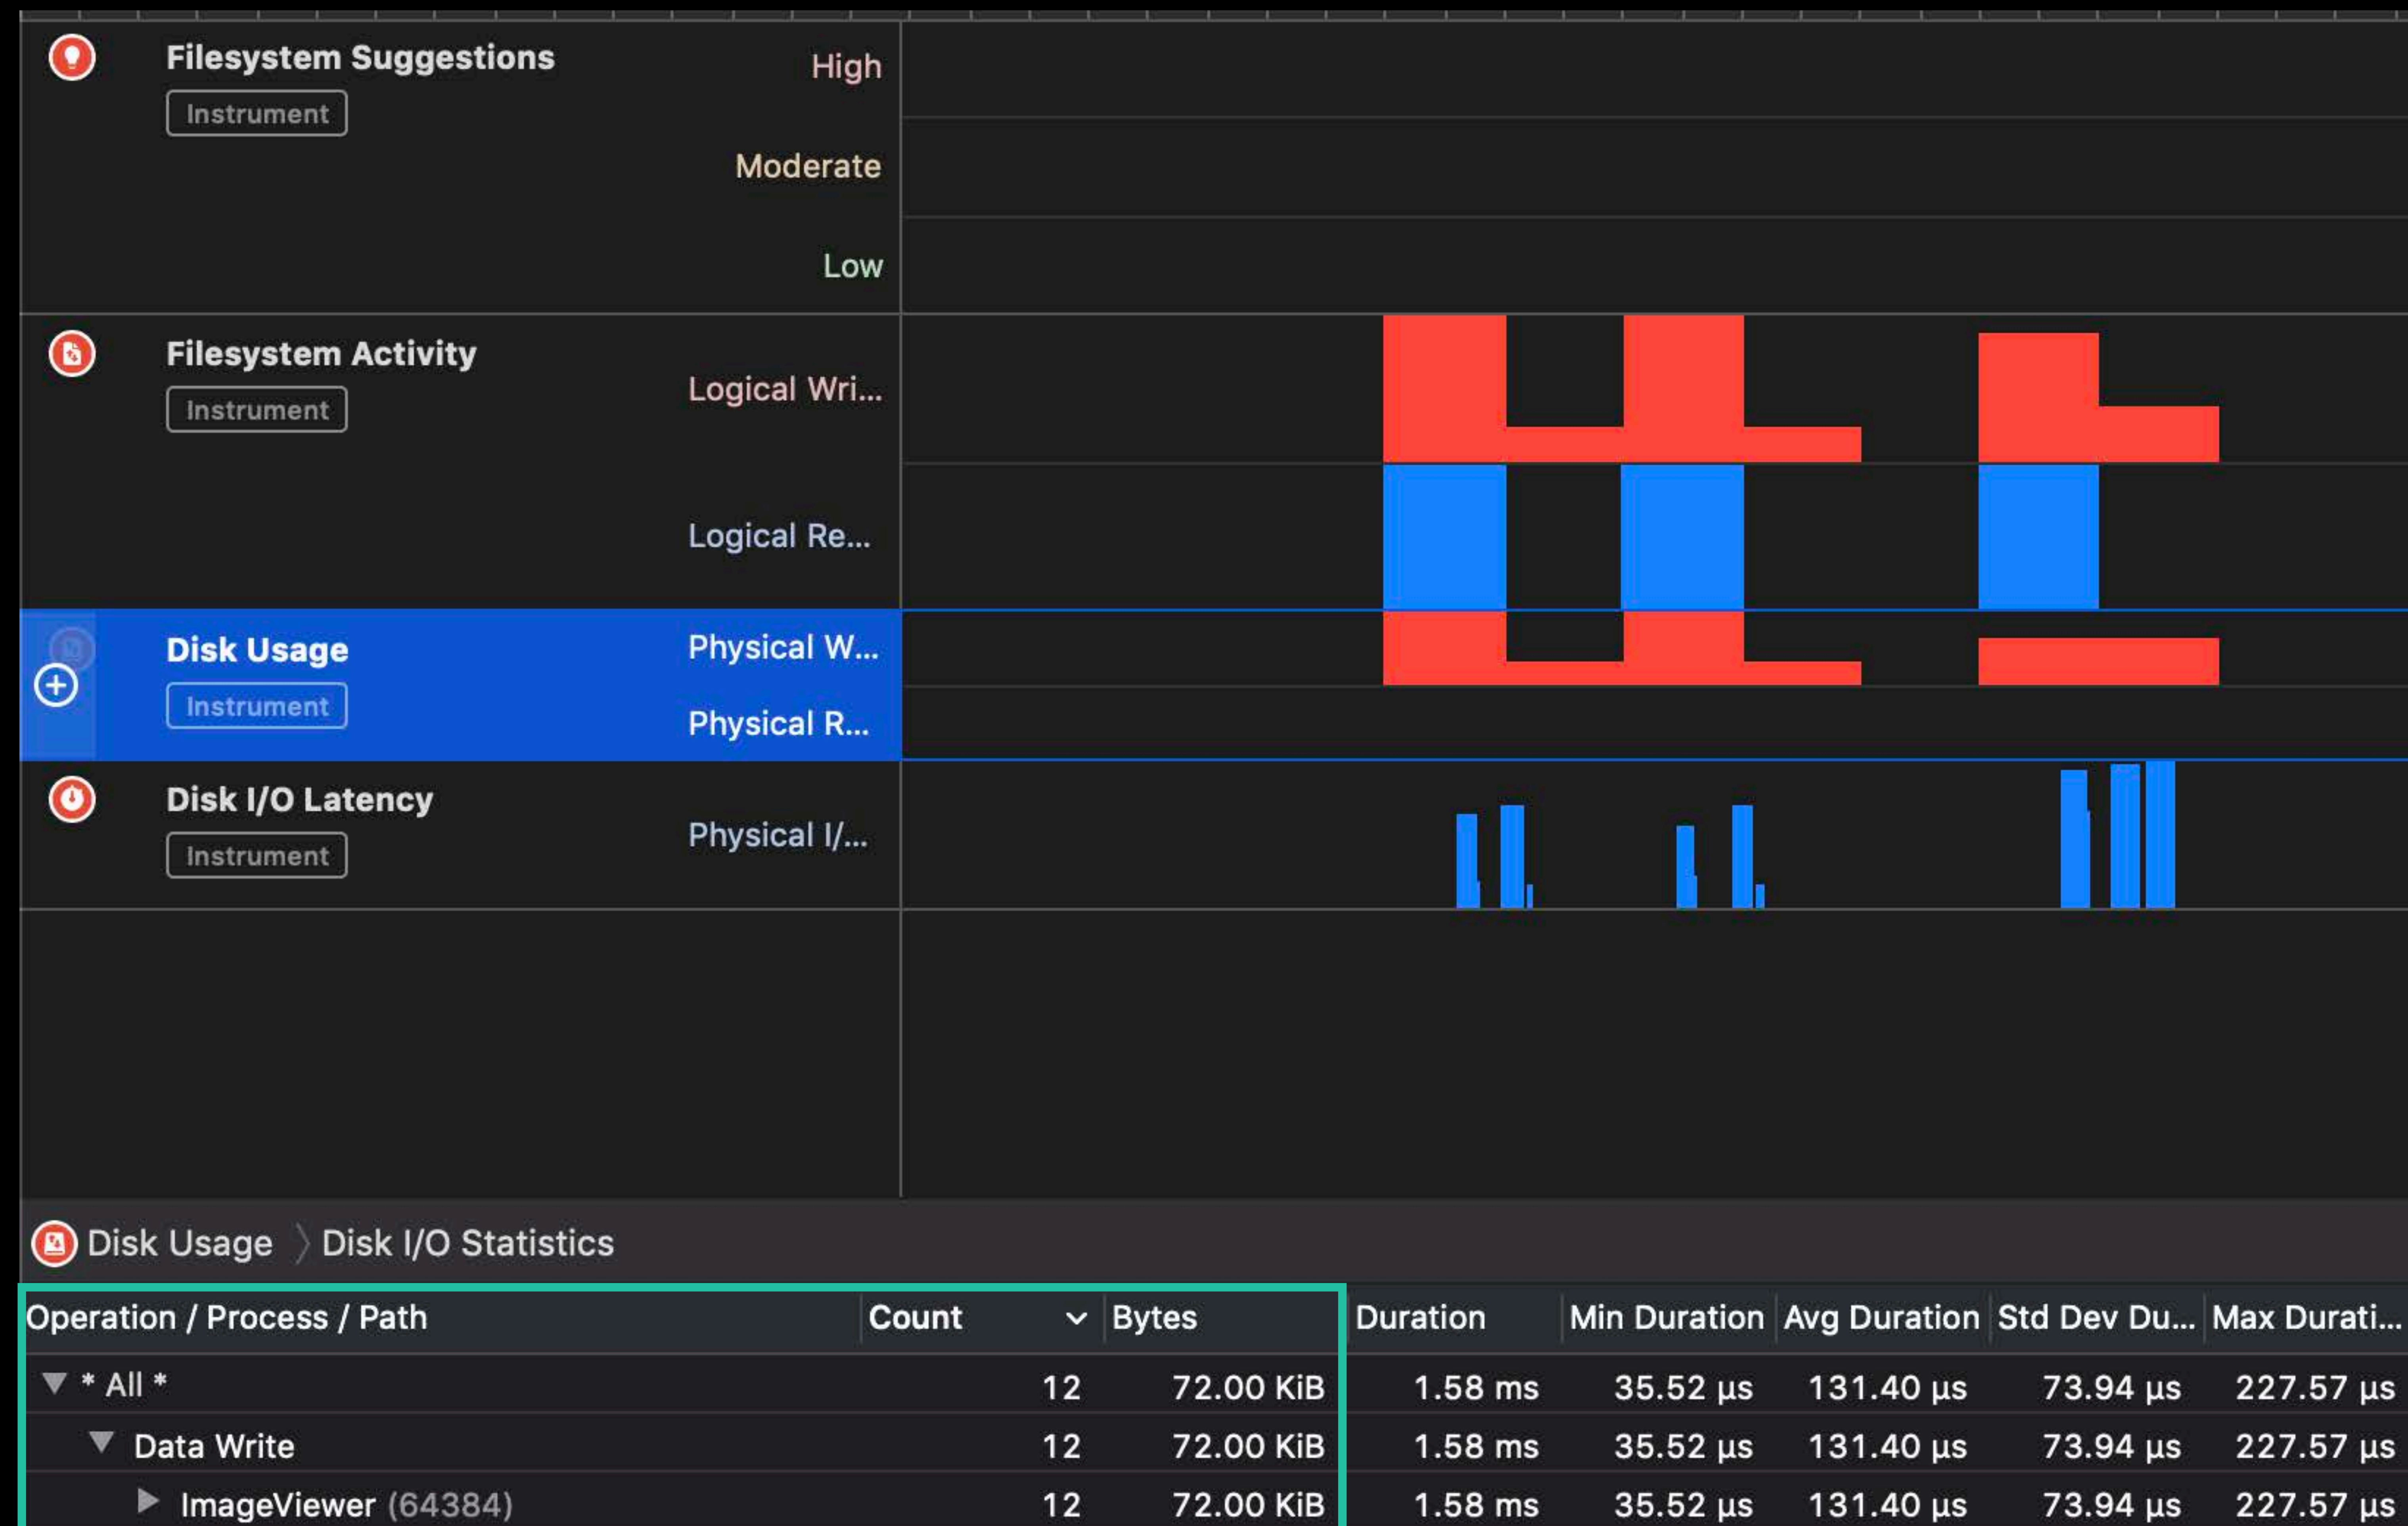


# Single Statement Transactions



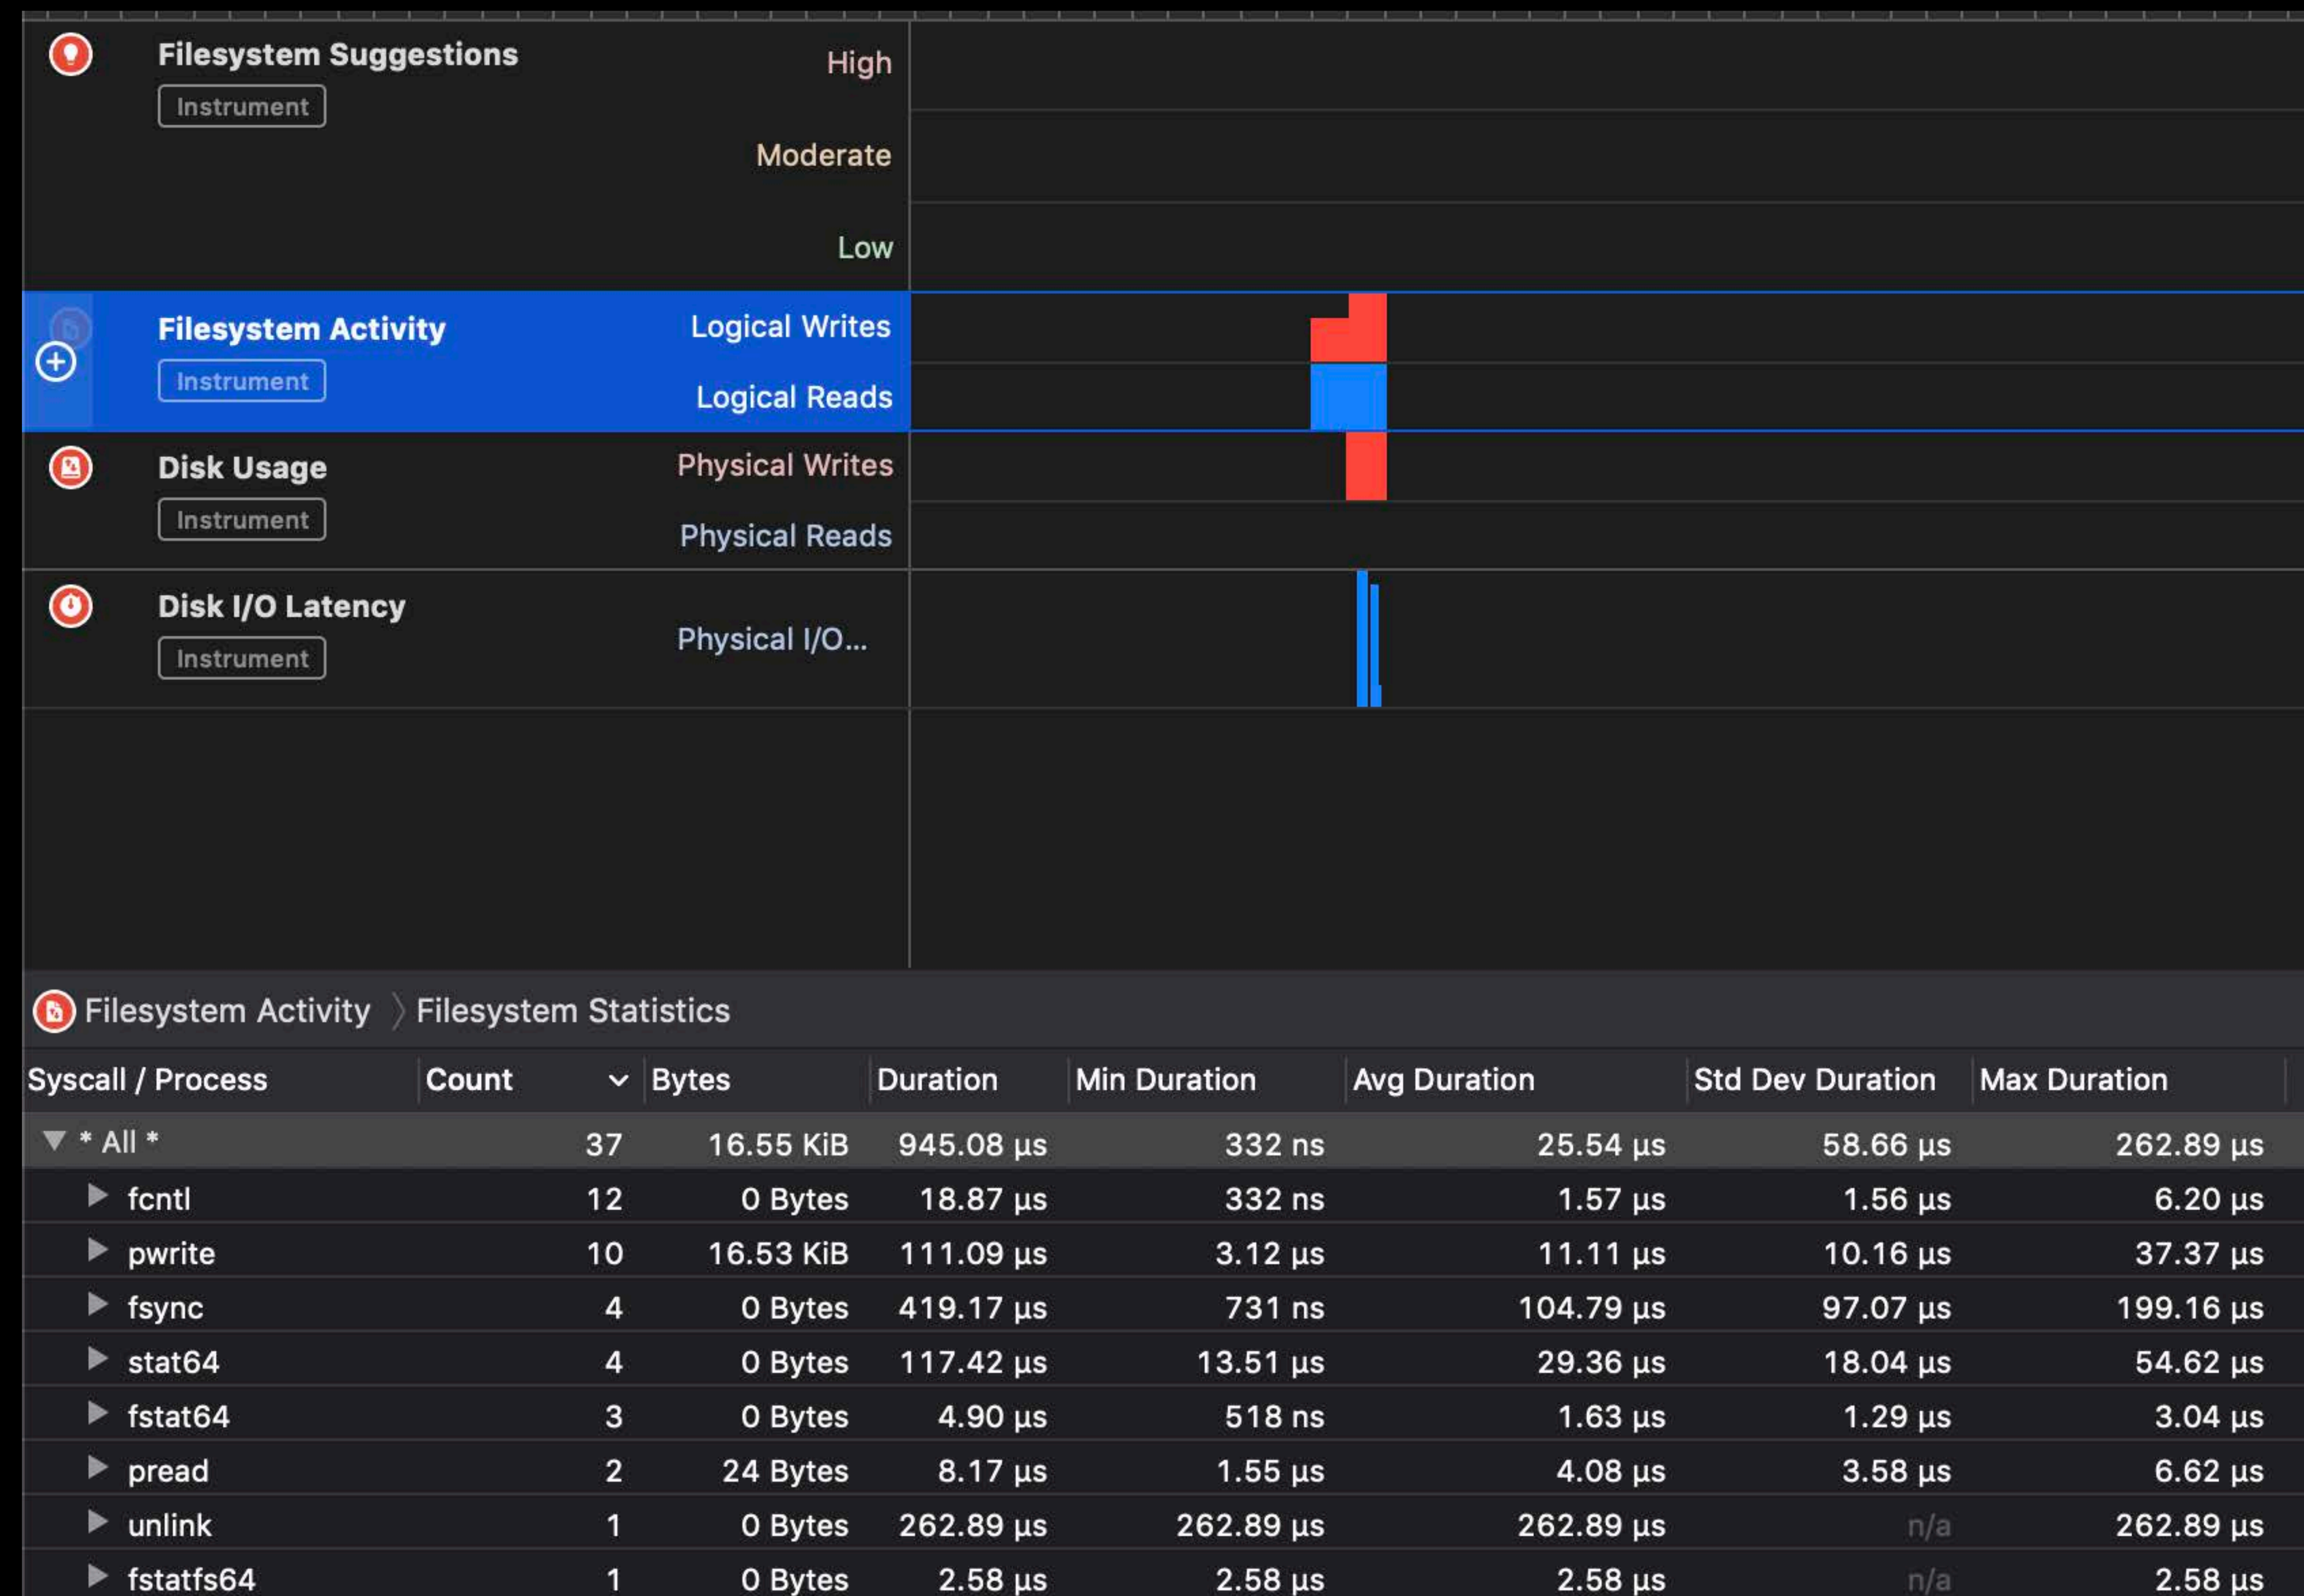


# Single Statement Transactions



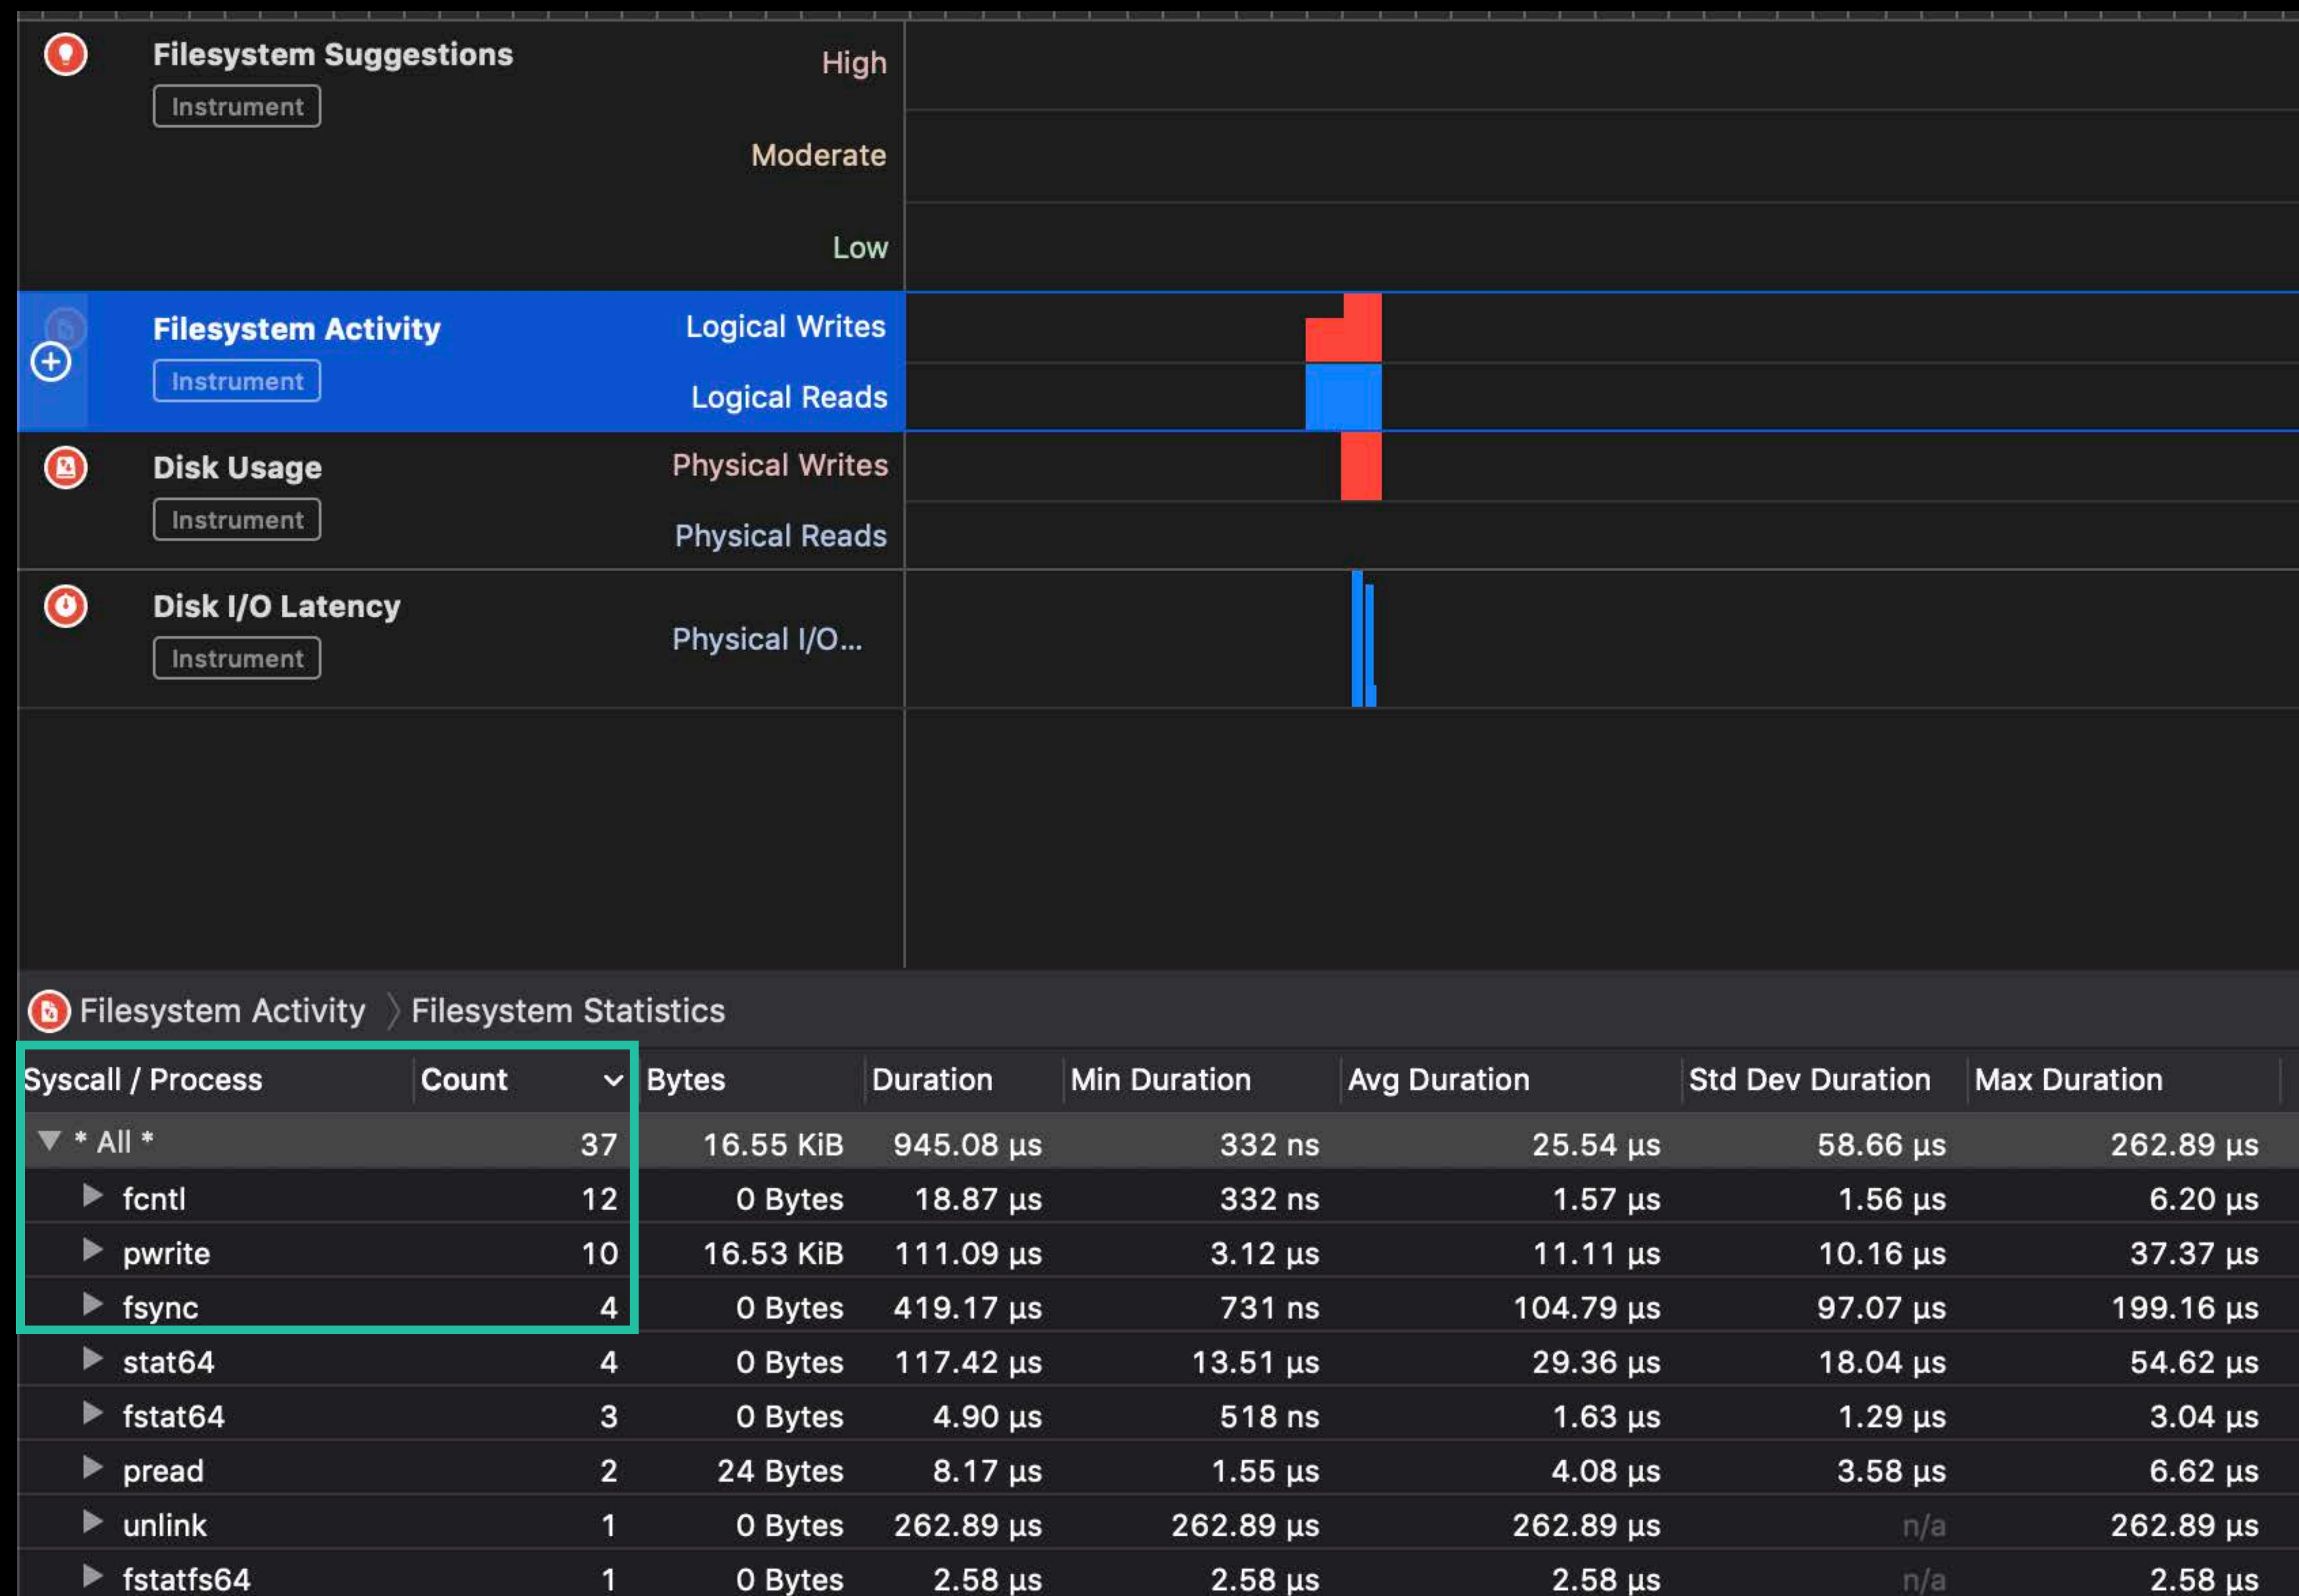


# Multiple Statement Transaction





# Multiple Statement Transaction





# Multiple Statement Transaction





# Multiple Statement Transaction

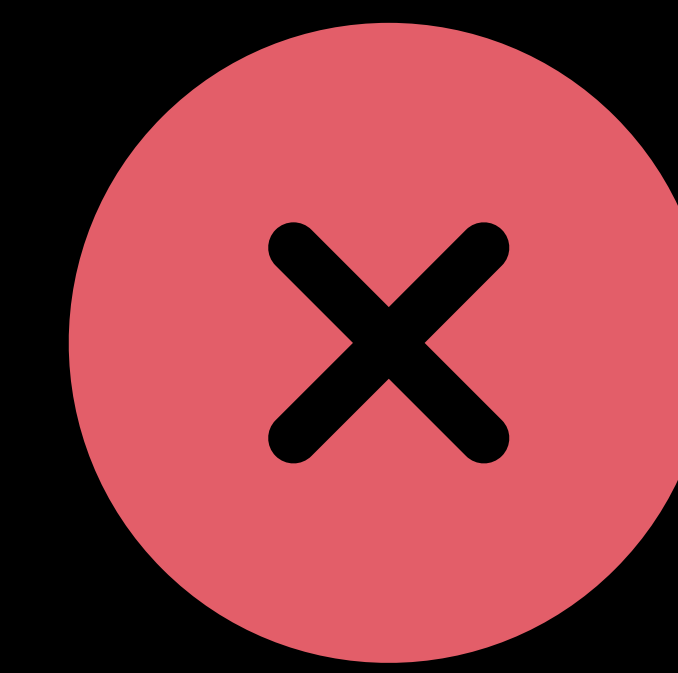




# Transactions Comparison



Multiple Statement



Single Statement

Filesystem Activity > Filesystem Statistics	
Syscall / Process	Count
▼ * All *	37
▶ fcntl	12
▶ pwrite	10
▶ fsync	4
▶ stat64	4
▶ fstat64	3
▶ pread	2
▶ unlink	1
▶ fstatfs64	1

Filesystem Activity > Filesystem Statistics	
Syscall / Process	Count
▼ * All *	111
▶ fcntl	36
▶ pwrite	30
▶ fsync	12
▶ stat64	12
▶ fstat64	9
▶ pread	6
▶ fstatfs64	3
▶ unlink	3

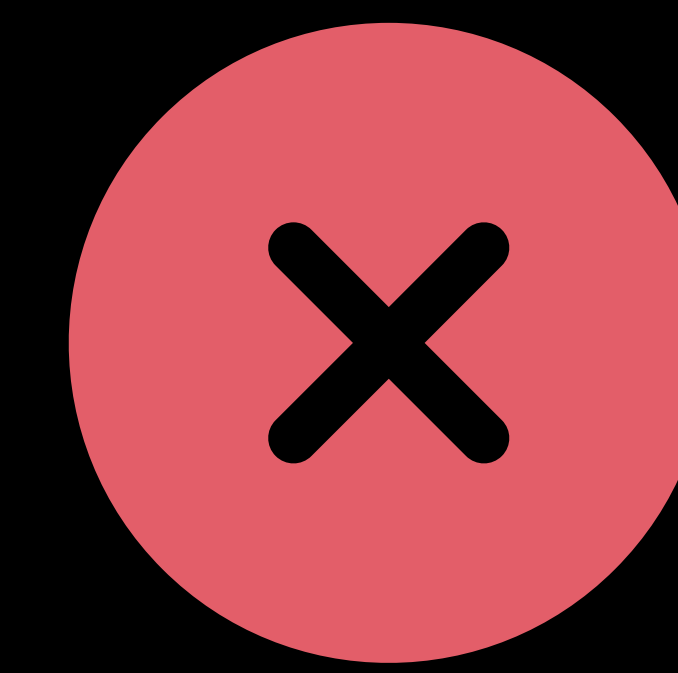


# Transactions Comparison



Multiple Statement

Disk Usage > Disk I/O Statistics			
Operation / Process / Path	Count	↓	Bytes
▼ * All *	4		24.00 KiB
▶ Data Write	4		24.00 KiB




Single Statement

Disk Usage > Disk I/O Statistics			
Operation / Process / Path	Count	↓	Bytes
▼ * All *	12		72.00 KiB
▼ Data Write →	12		72.00 KiB
▶ ImageViewer (64384)	12		72.00 KiB




# Full Vacuum

 Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	27	168.00 KiB	1.16 ms	18.91 µs	43.12 µs	57.25 µs	276.55 µs
▼ Data Write	27	168.00 KiB	1.16 ms	18.91 µs	43.12 µs	57.25 µs	276.55 µs
▼ ImageViewer (65819)	27	168.00 KiB	1.16 ms	18.91 µs	43.12 µs	57.25 µs	276.55 µs
/ImageViewer/full_vacuum.db-journal	18	108.00 KiB	932.60 µs	18.93 µs	51.81 µs	68.99 µs	276.55 µs
/ImageViewer/full_vacuum.db	9	60.00 KiB	231.75 µs	18.91 µs	25.75 µs	5.52 µs	35.35 µs




# Full Vacuum

 Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	27	168.00 KiB	1.16 ms	18.91 µs	43.12 µs	57.25 µs	276.55 µs
▼ Data Write	27	168.00 KiB	1.16 ms	18.91 µs	43.12 µs	57.25 µs	276.55 µs
▼ ImageViewer (65819)	27	168.00 KiB	1.16 ms	18.91 µs	43.12 µs	57.25 µs	276.55 µs
/ImageViewer/full_vacuum.db-journal	18	108.00 KiB	932.60 µs	18.93 µs	51.81 µs	68.99 µs	276.55 µs
/ImageViewer/full_vacuum.db	9	60.00 KiB	231.75 µs	18.91 µs	25.75 µs	5.52 µs	35.35 µs




# Incremental Vacuum

 Disk Usage > Disk I/O Statistics

Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	12	72.00 KiB	882.60 μs	23.22 μs	73.55 μs	88.51 μs	301.52 μs
▼ Data Write	12	72.00 KiB	882.60 μs	23.22 μs	73.55 μs	88.51 μs	301.52 μs
▼ ImageViewer (66075)	12	72.00 KiB	882.60 μs	23.22 μs	73.55 μs	88.51 μs	301.52 μs
/ImageViewer/incremental_vacuum.db-journal	9	48.00 KiB	785.72 μs	23.22 μs	87.30 μs	99.60 μs	301.52 μs
/ImageViewer/incremental_vacuum.db	3	24.00 KiB	96.88 μs	31.62 μs	32.29 μs	805 ns	33.18 μs



# Incremental Vacuum

 Disk Usage > Disk I/O Statistics

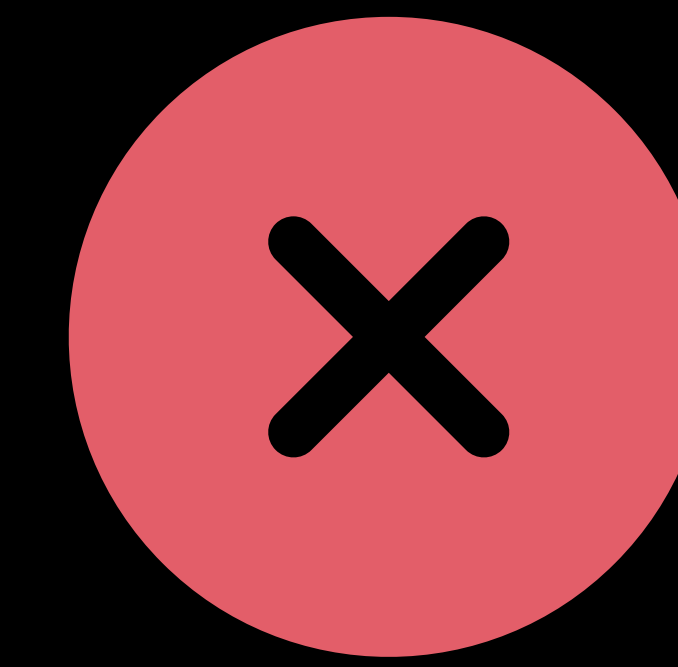
Operation / Process / Path	Count	Bytes	Total Laten...	Min Total L...	Avg Total L...	Std Dev To...	Max Total L...
▼ * All *	12	72.00 KiB	882.60 μs	23.22 μs	73.55 μs	88.51 μs	301.52 μs
▼ Data Write	12	72.00 KiB	882.60 μs	23.22 μs	73.55 μs	88.51 μs	301.52 μs
▼ ImageViewer (66075)	12	72.00 KiB	882.60 μs	23.22 μs	73.55 μs	88.51 μs	301.52 μs
/ImageViewer/incremental_vacuum.db-journal	9	48.00 KiB	785.72 μs	23.22 μs	87.30 μs	99.60 μs	301.52 μs
/ImageViewer/incremental_vacuum.db	3	24.00 KiB	96.88 μs	31.62 μs	32.29 μs	805 ns	33.18 μs



# Vacuuming Comparison



Incremental



Full

Disk Usage > Disk I/O Statistics			
Operation / Process / Path	Count	↓	Bytes
▼ * All *	12		72.00 KiB
▼ Data Write	12		72.00 KiB
▼ ImageViewer (66075)	12		72.00 KiB
/ImageViewer/incremental_vacuum.db-journal	9		48.00 KiB
/ImageViewer/incremental_vacuum.db	3		24.00 KiB

Disk Usage > Disk I/O Statistics			
Operation / Process / Path	Count	↓	Bytes
▼ * All *	27		168.00 KiB
▼ Data Write	27		168.00 KiB
▼ ImageViewer (65819)	27		168.00 KiB
/ImageViewer/full_vacuum.db-journal	18		108.00 KiB
/ImageViewer/full_vacuum.db	9		60.00 KiB

# Summary



# Summary

Apply these lessons

# Summary

Apply these lessons

Profile with File Activity Instrument



# Summary

Apply these lessons

Profile with File Activity Instrument

Continue optimizing storage!

# More Information

[developer.apple.com/wwdc19/419](https://developer.apple.com/wwdc19/419)

---

Making Apps with Core Data

WWDC 2019

---

Performance, Power, Crashes, and Debugging Lab

Friday, 3:00

---



