



Mac OS X File Systems



Had to give it a shot!

Most ideas come from previous ideas.

Alan Kay

**Looking
Back**





WITH NO BALLS

145 COMODOS
PAÑALES

PAÑAL
DESCARTABLE

100%
KIDS

Unix

BY DRYPERS

UNISEX

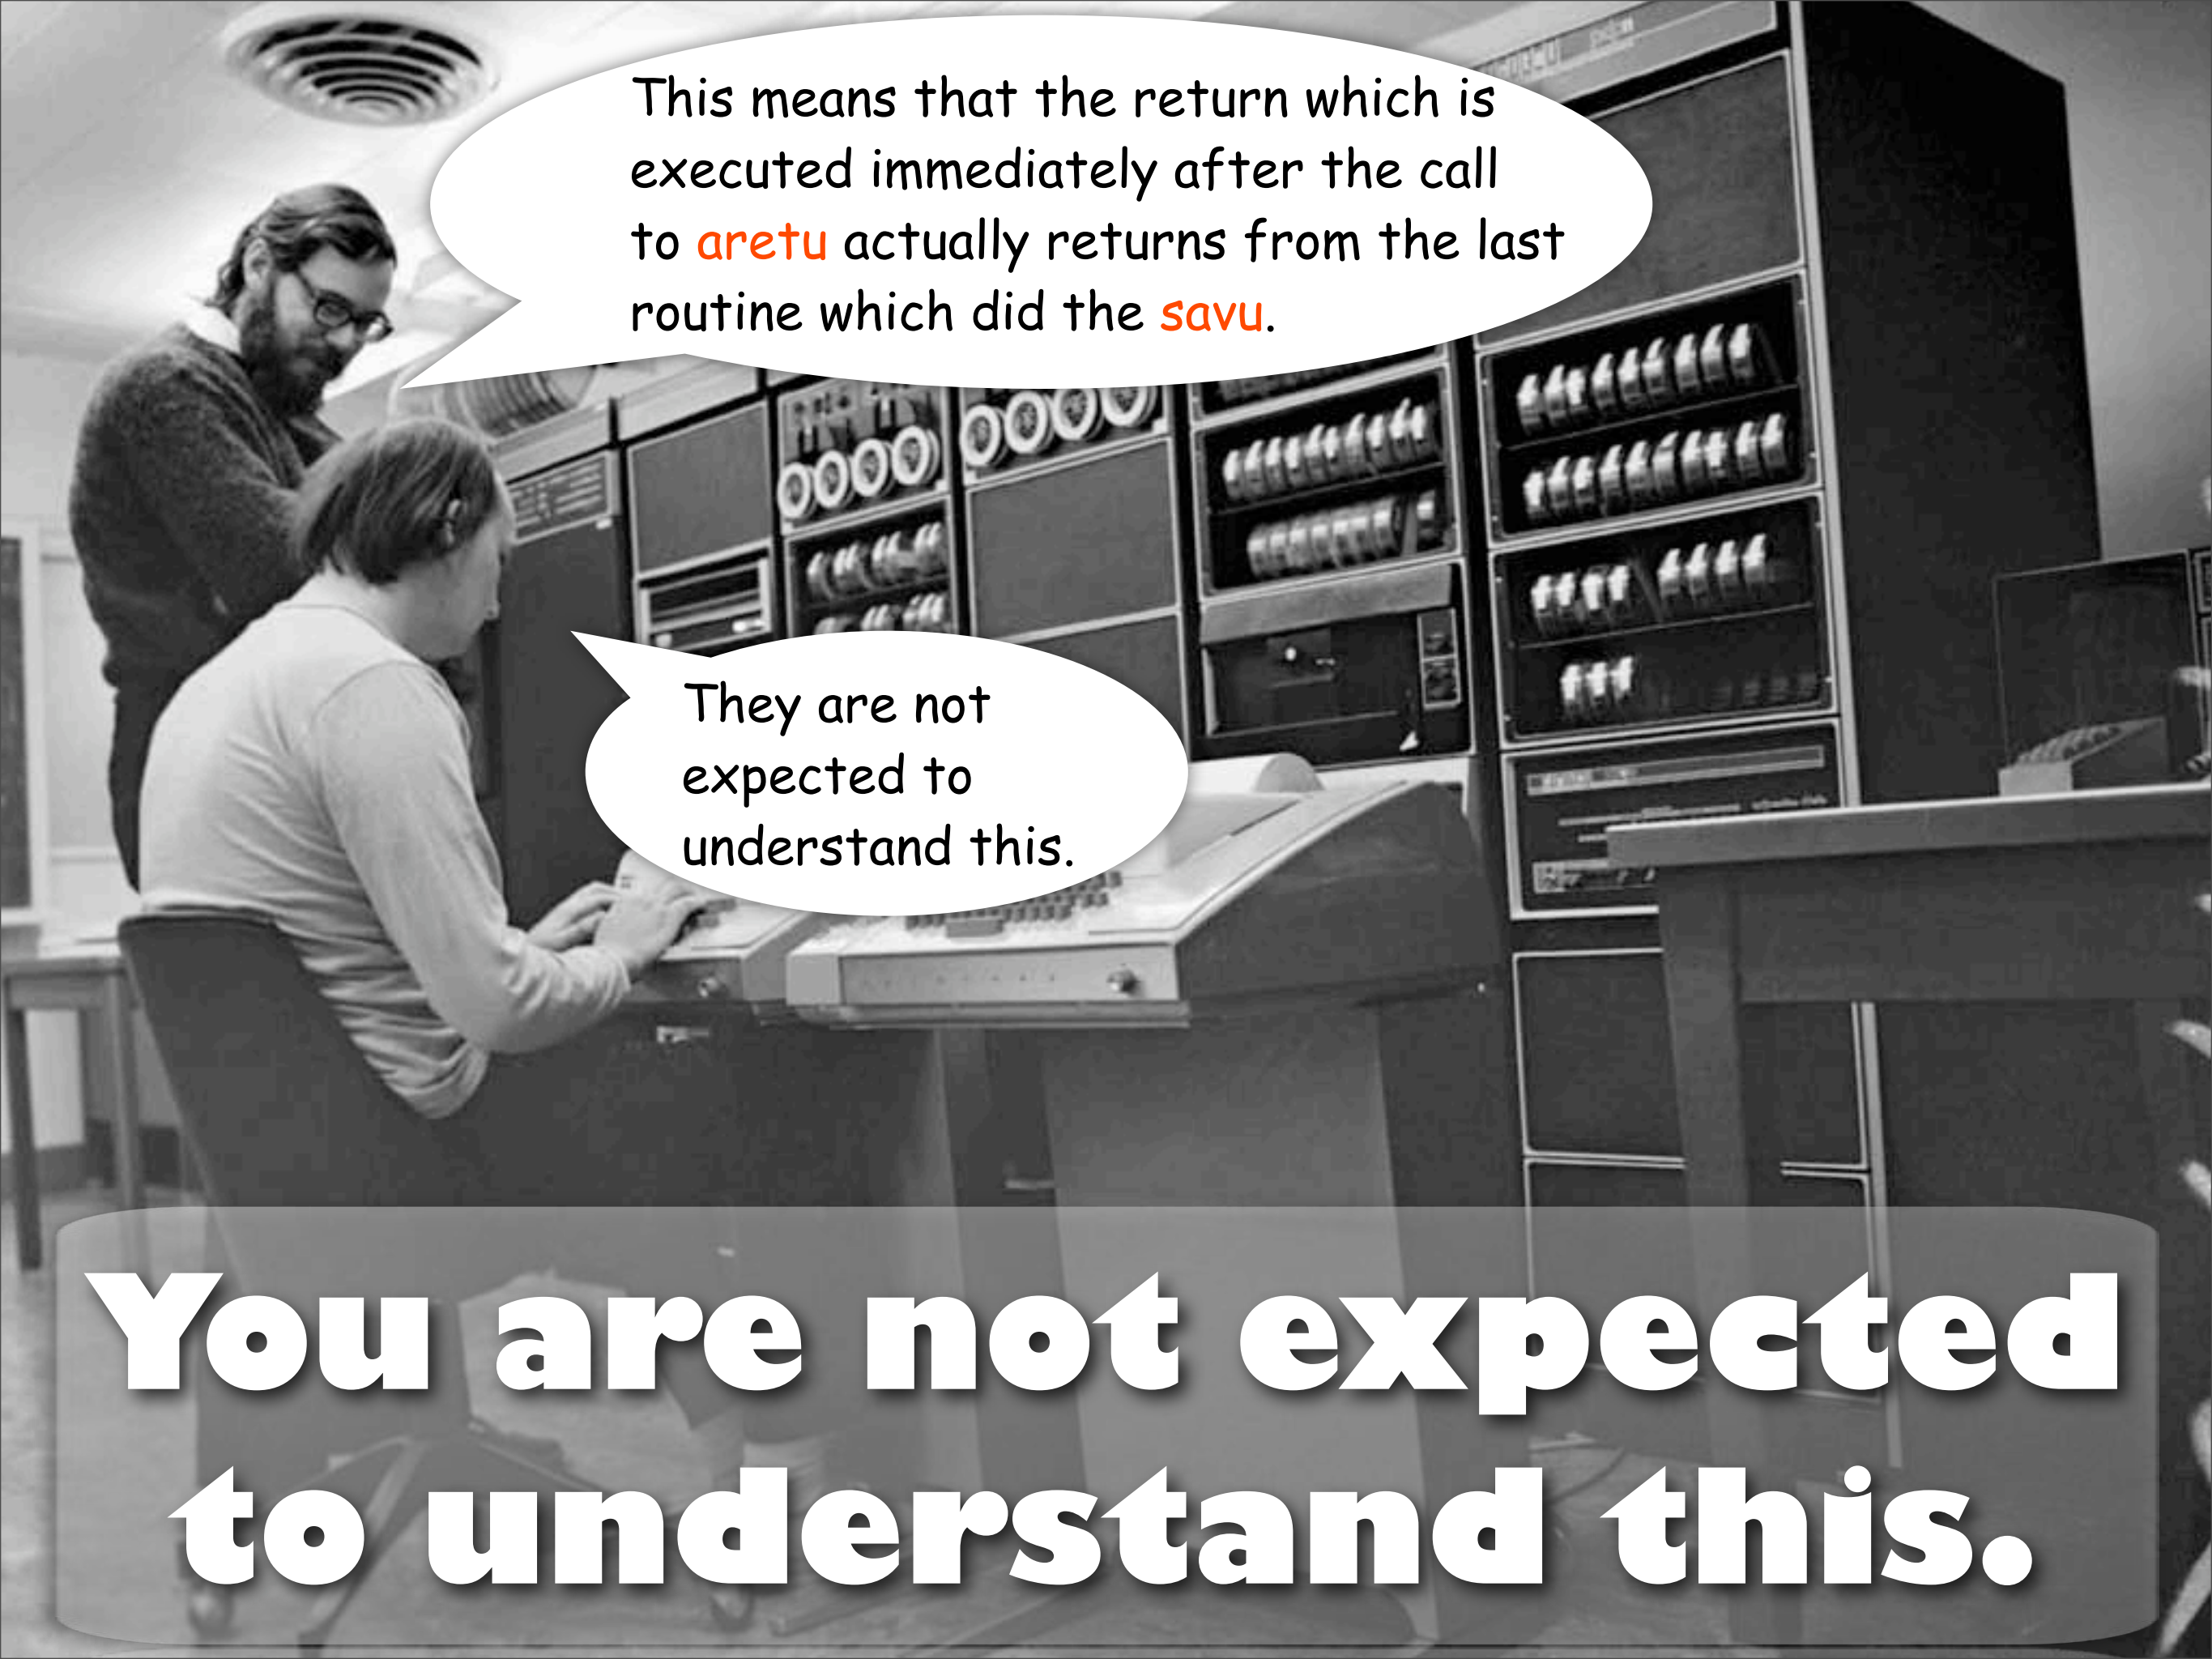
DOBLE ABSORCIÓN
DOBLE BARRERA
ULTRADELGADO



CONTENIDO

14

COMODIDAD



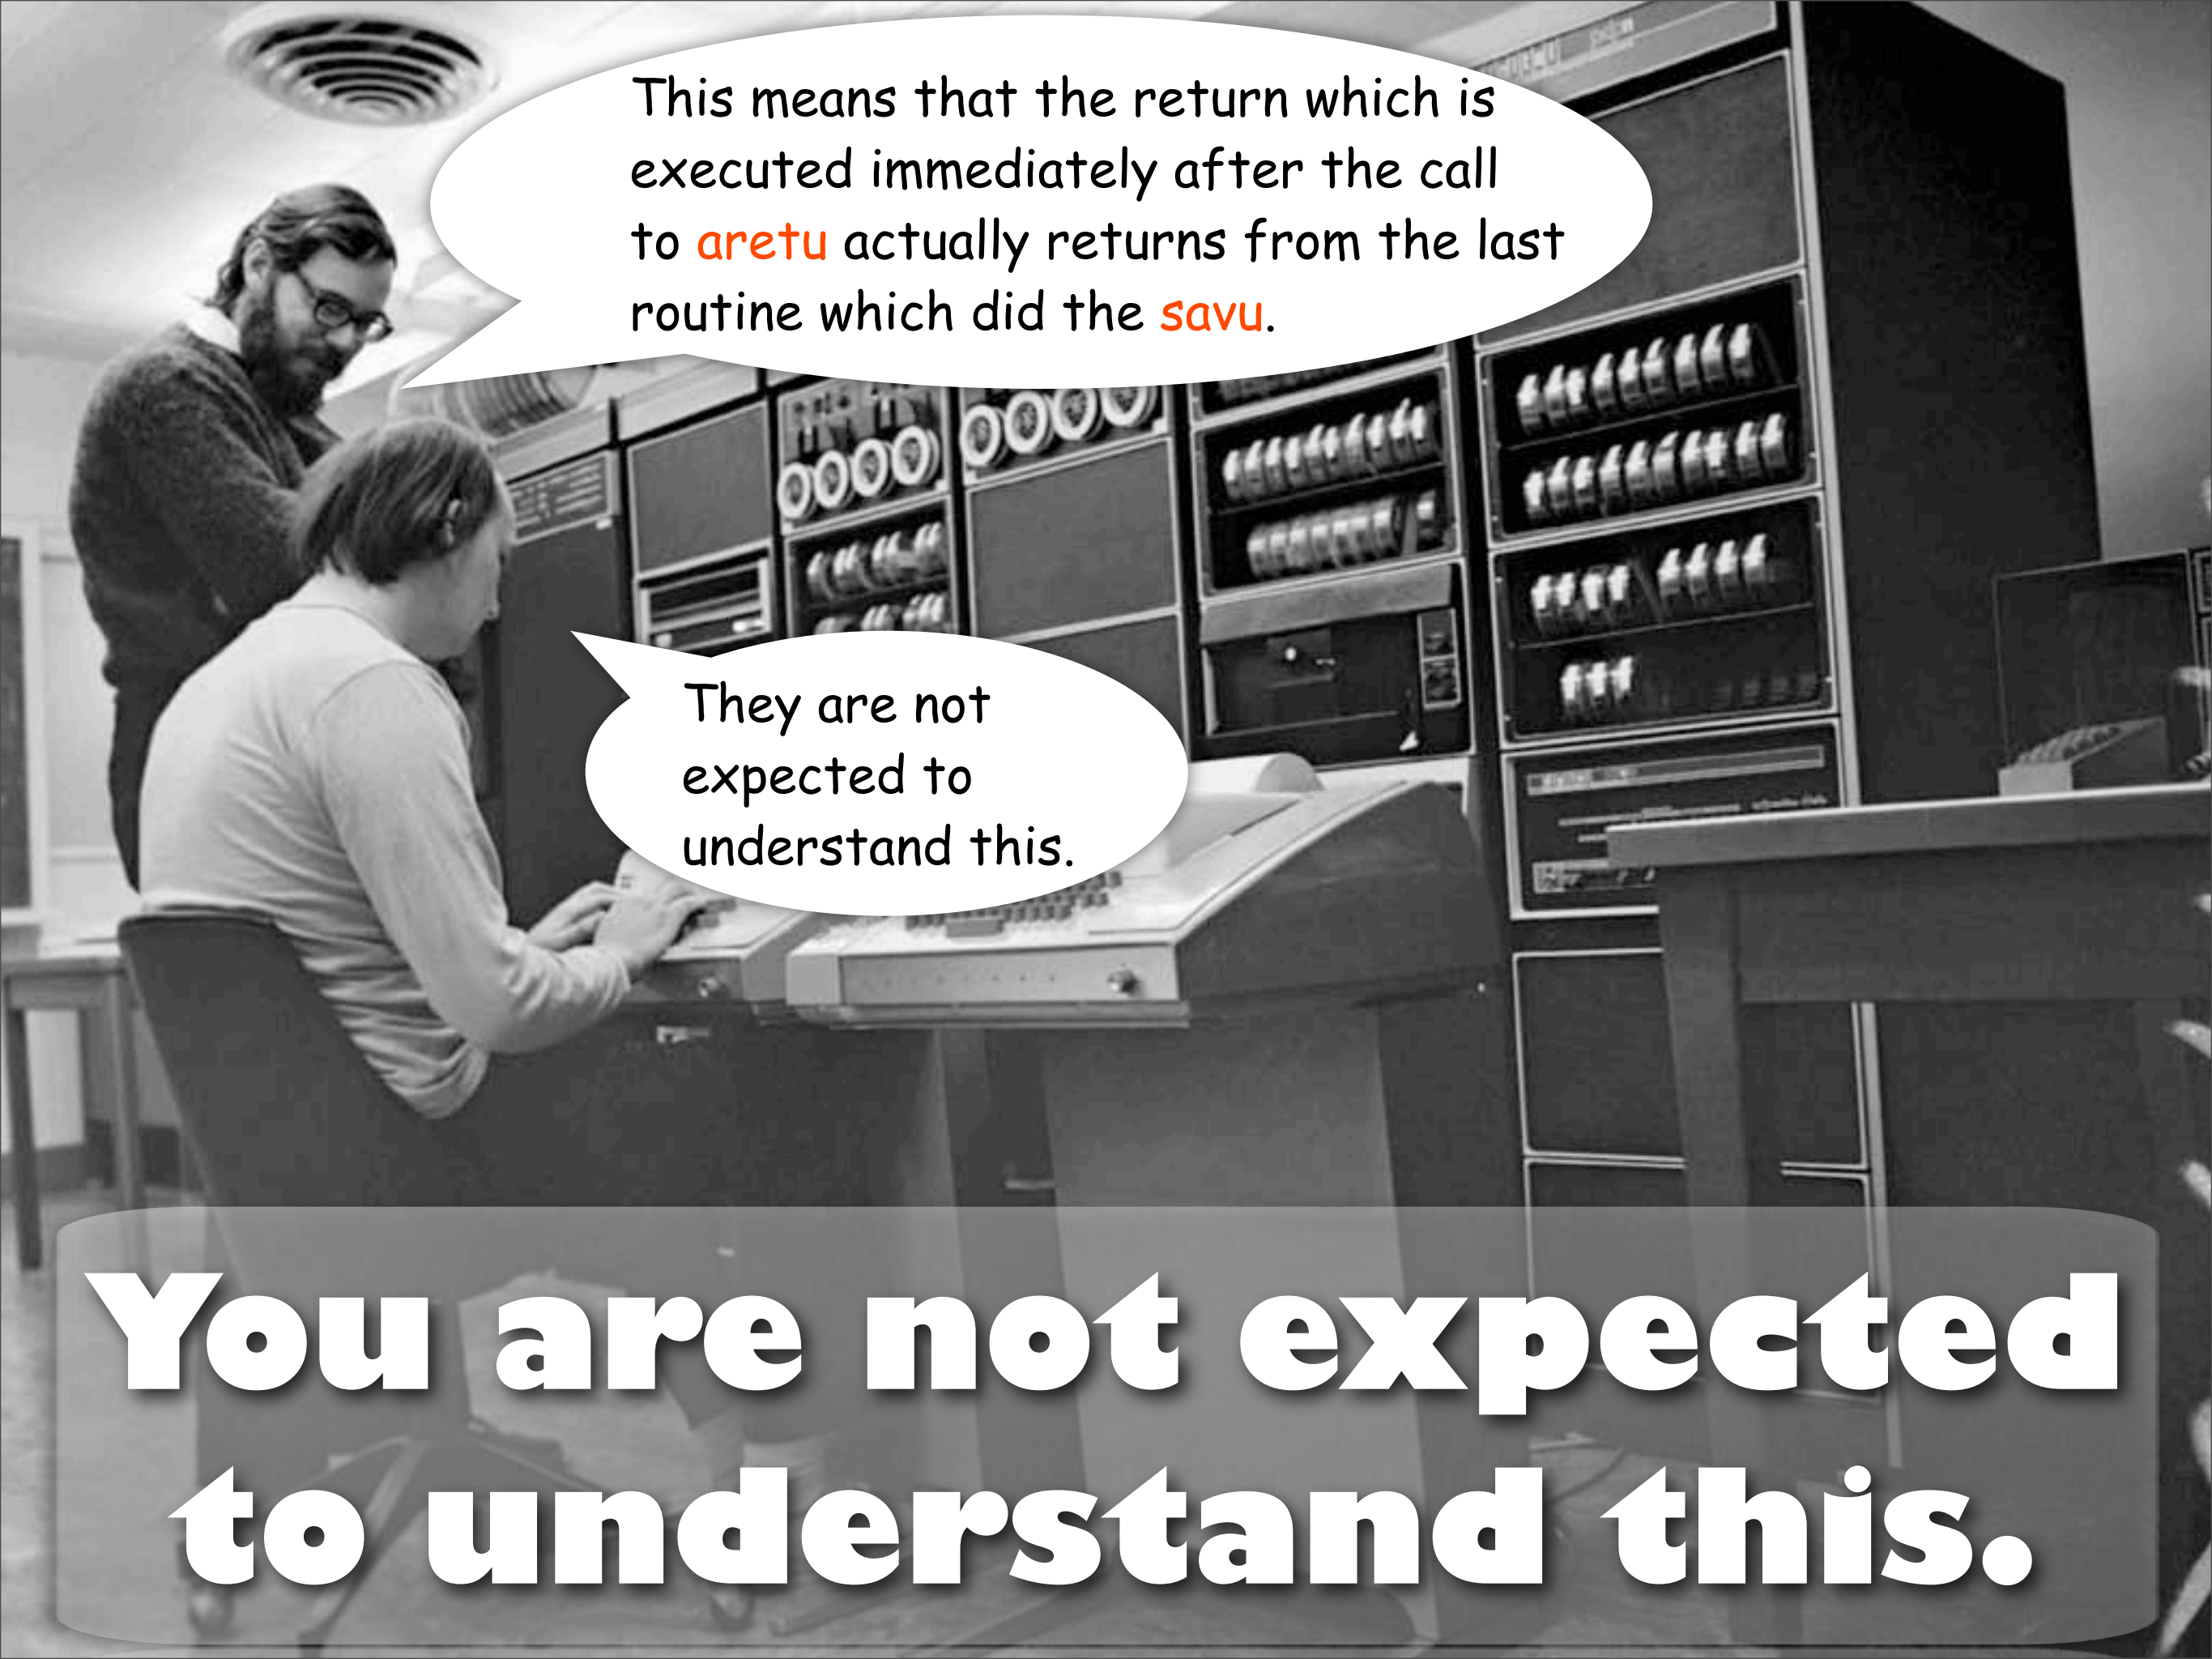
This means that the return which is executed immediately after the call to **aretu** actually returns from the last routine which did the **savu**.

They are not expected to understand this.

You are not expected to understand this.

UNIX

**The most important
role of UNIX is to
provide a file system.**



This means that the return which is executed immediately after the call to **aretu** actually returns from the last routine which did the **savu**.

They are not expected to understand this.

You are not expected to understand this.

UNIX

**The most important
Everything is a file.
provide a file system.**

User Program

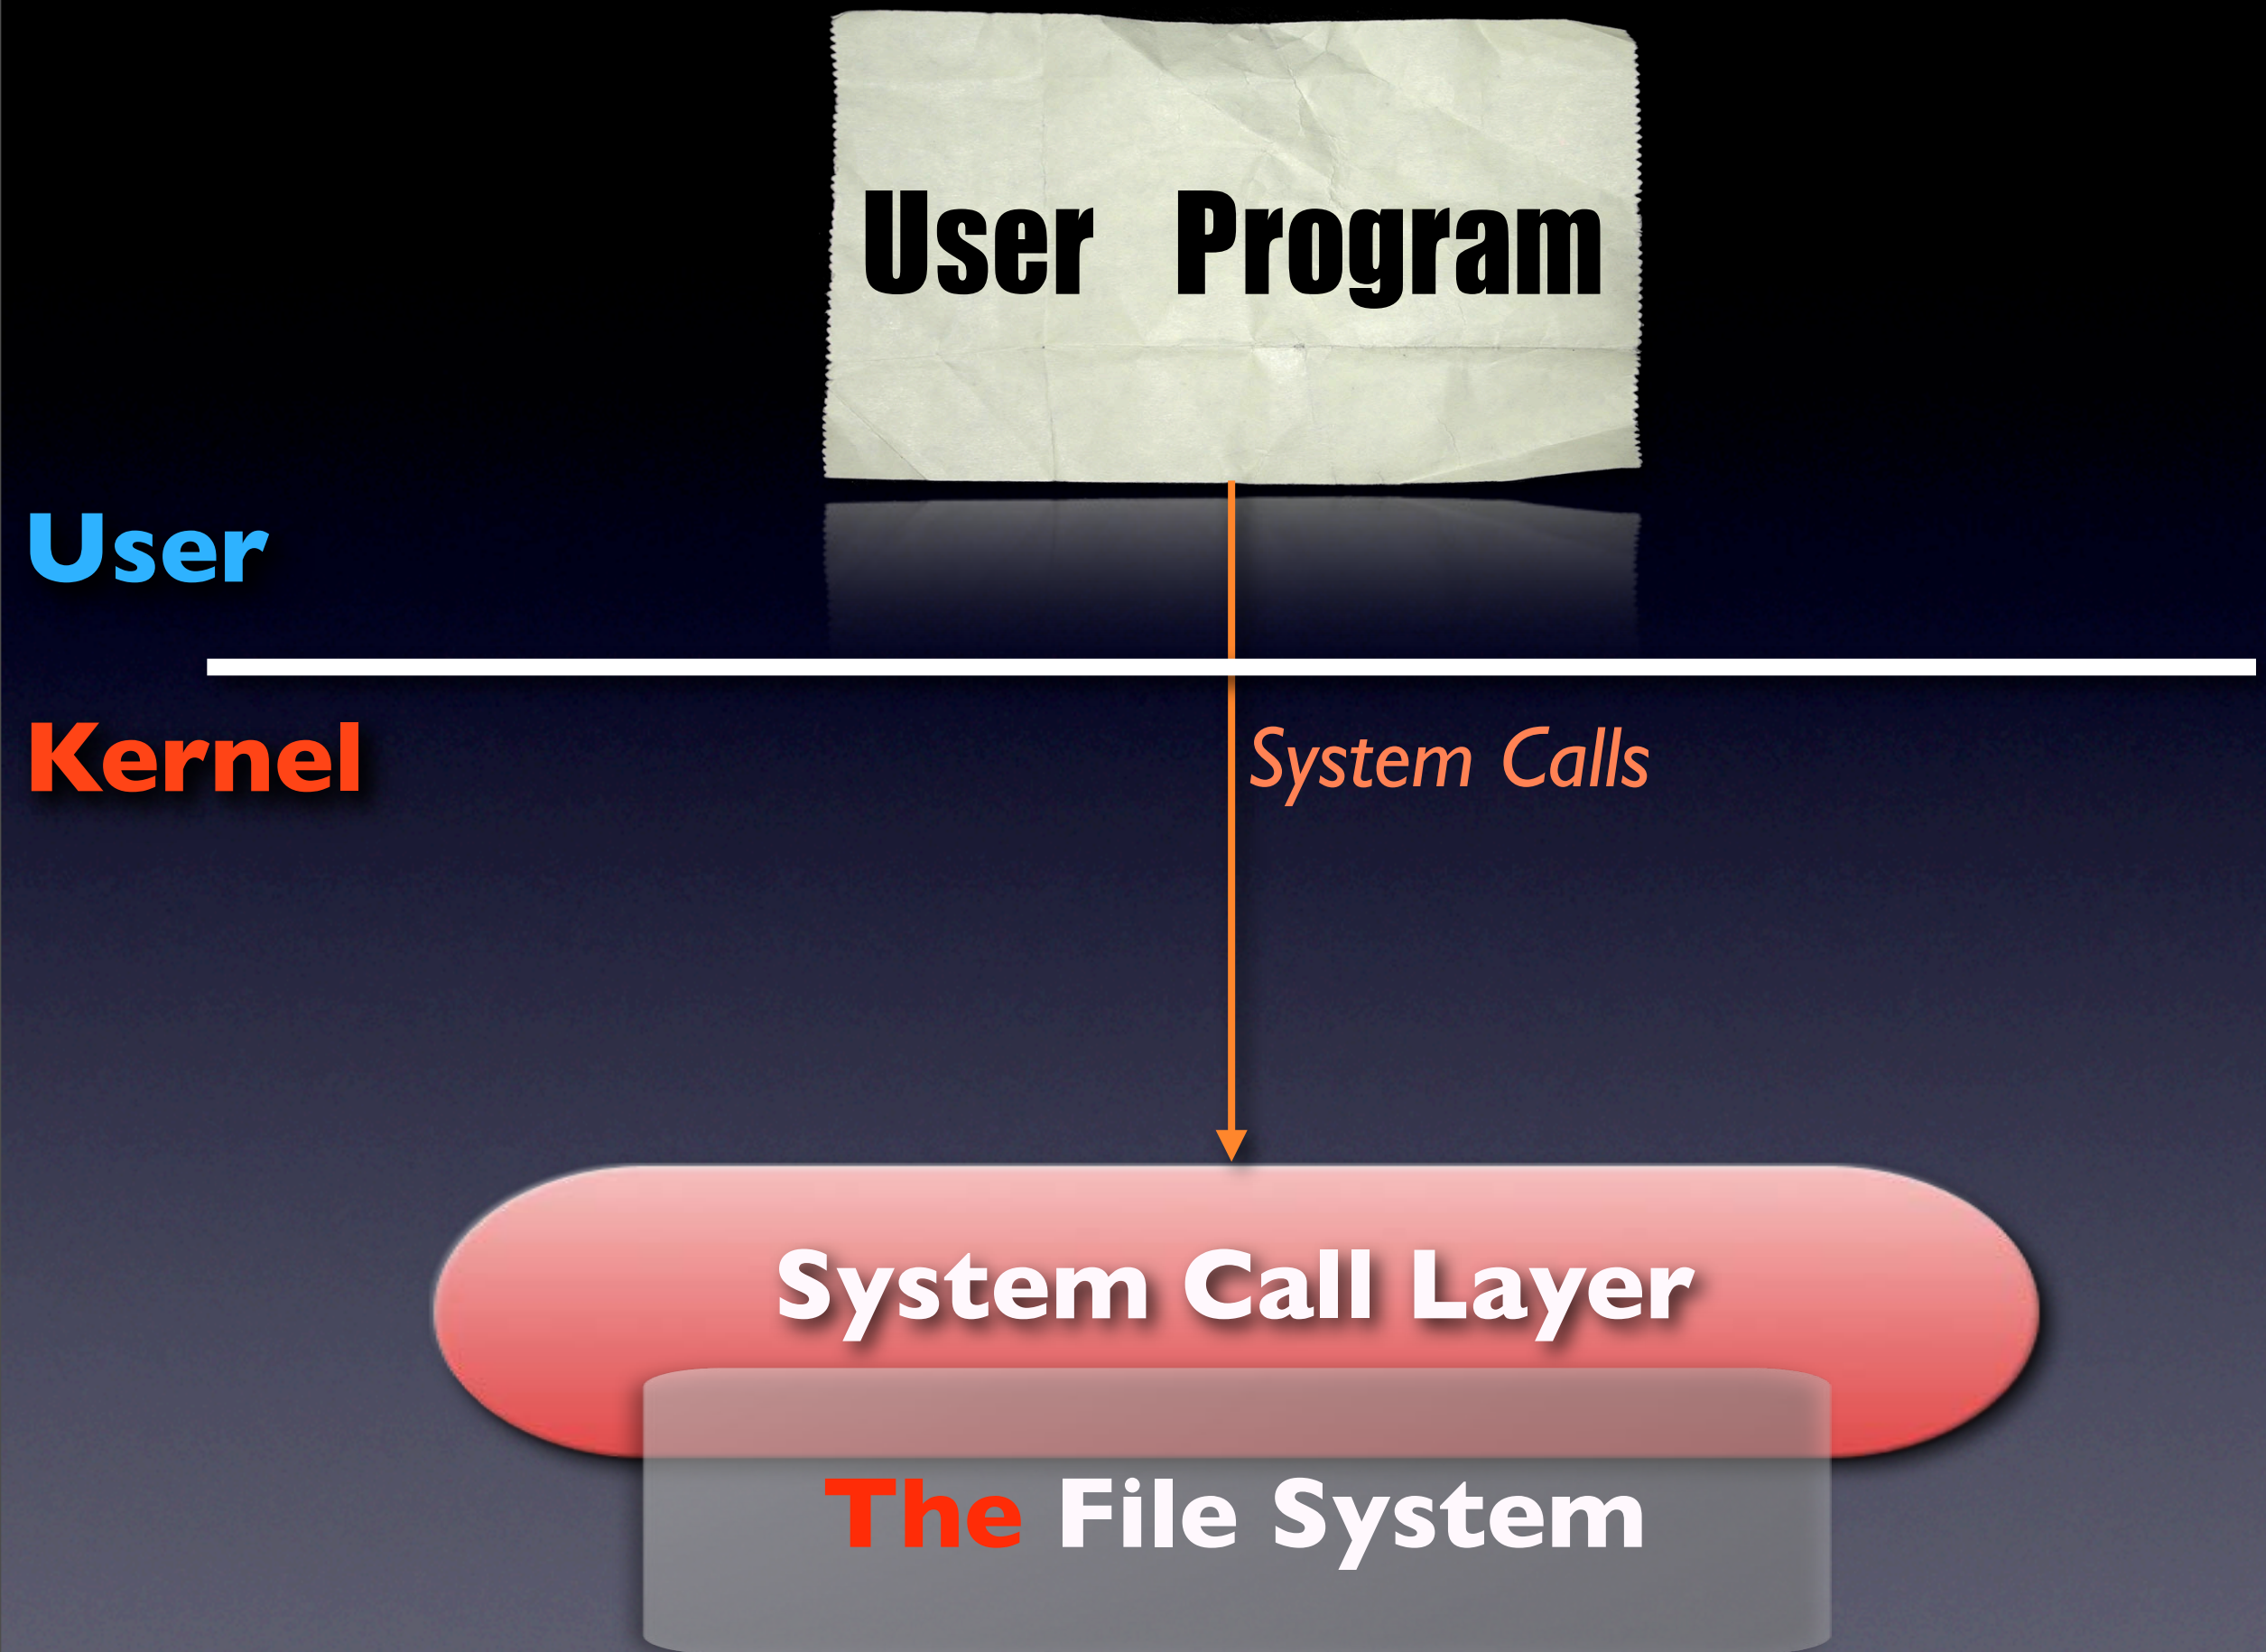
User

Kernel

System Calls

System Call Layer

The File System



User Program

User

Kernel

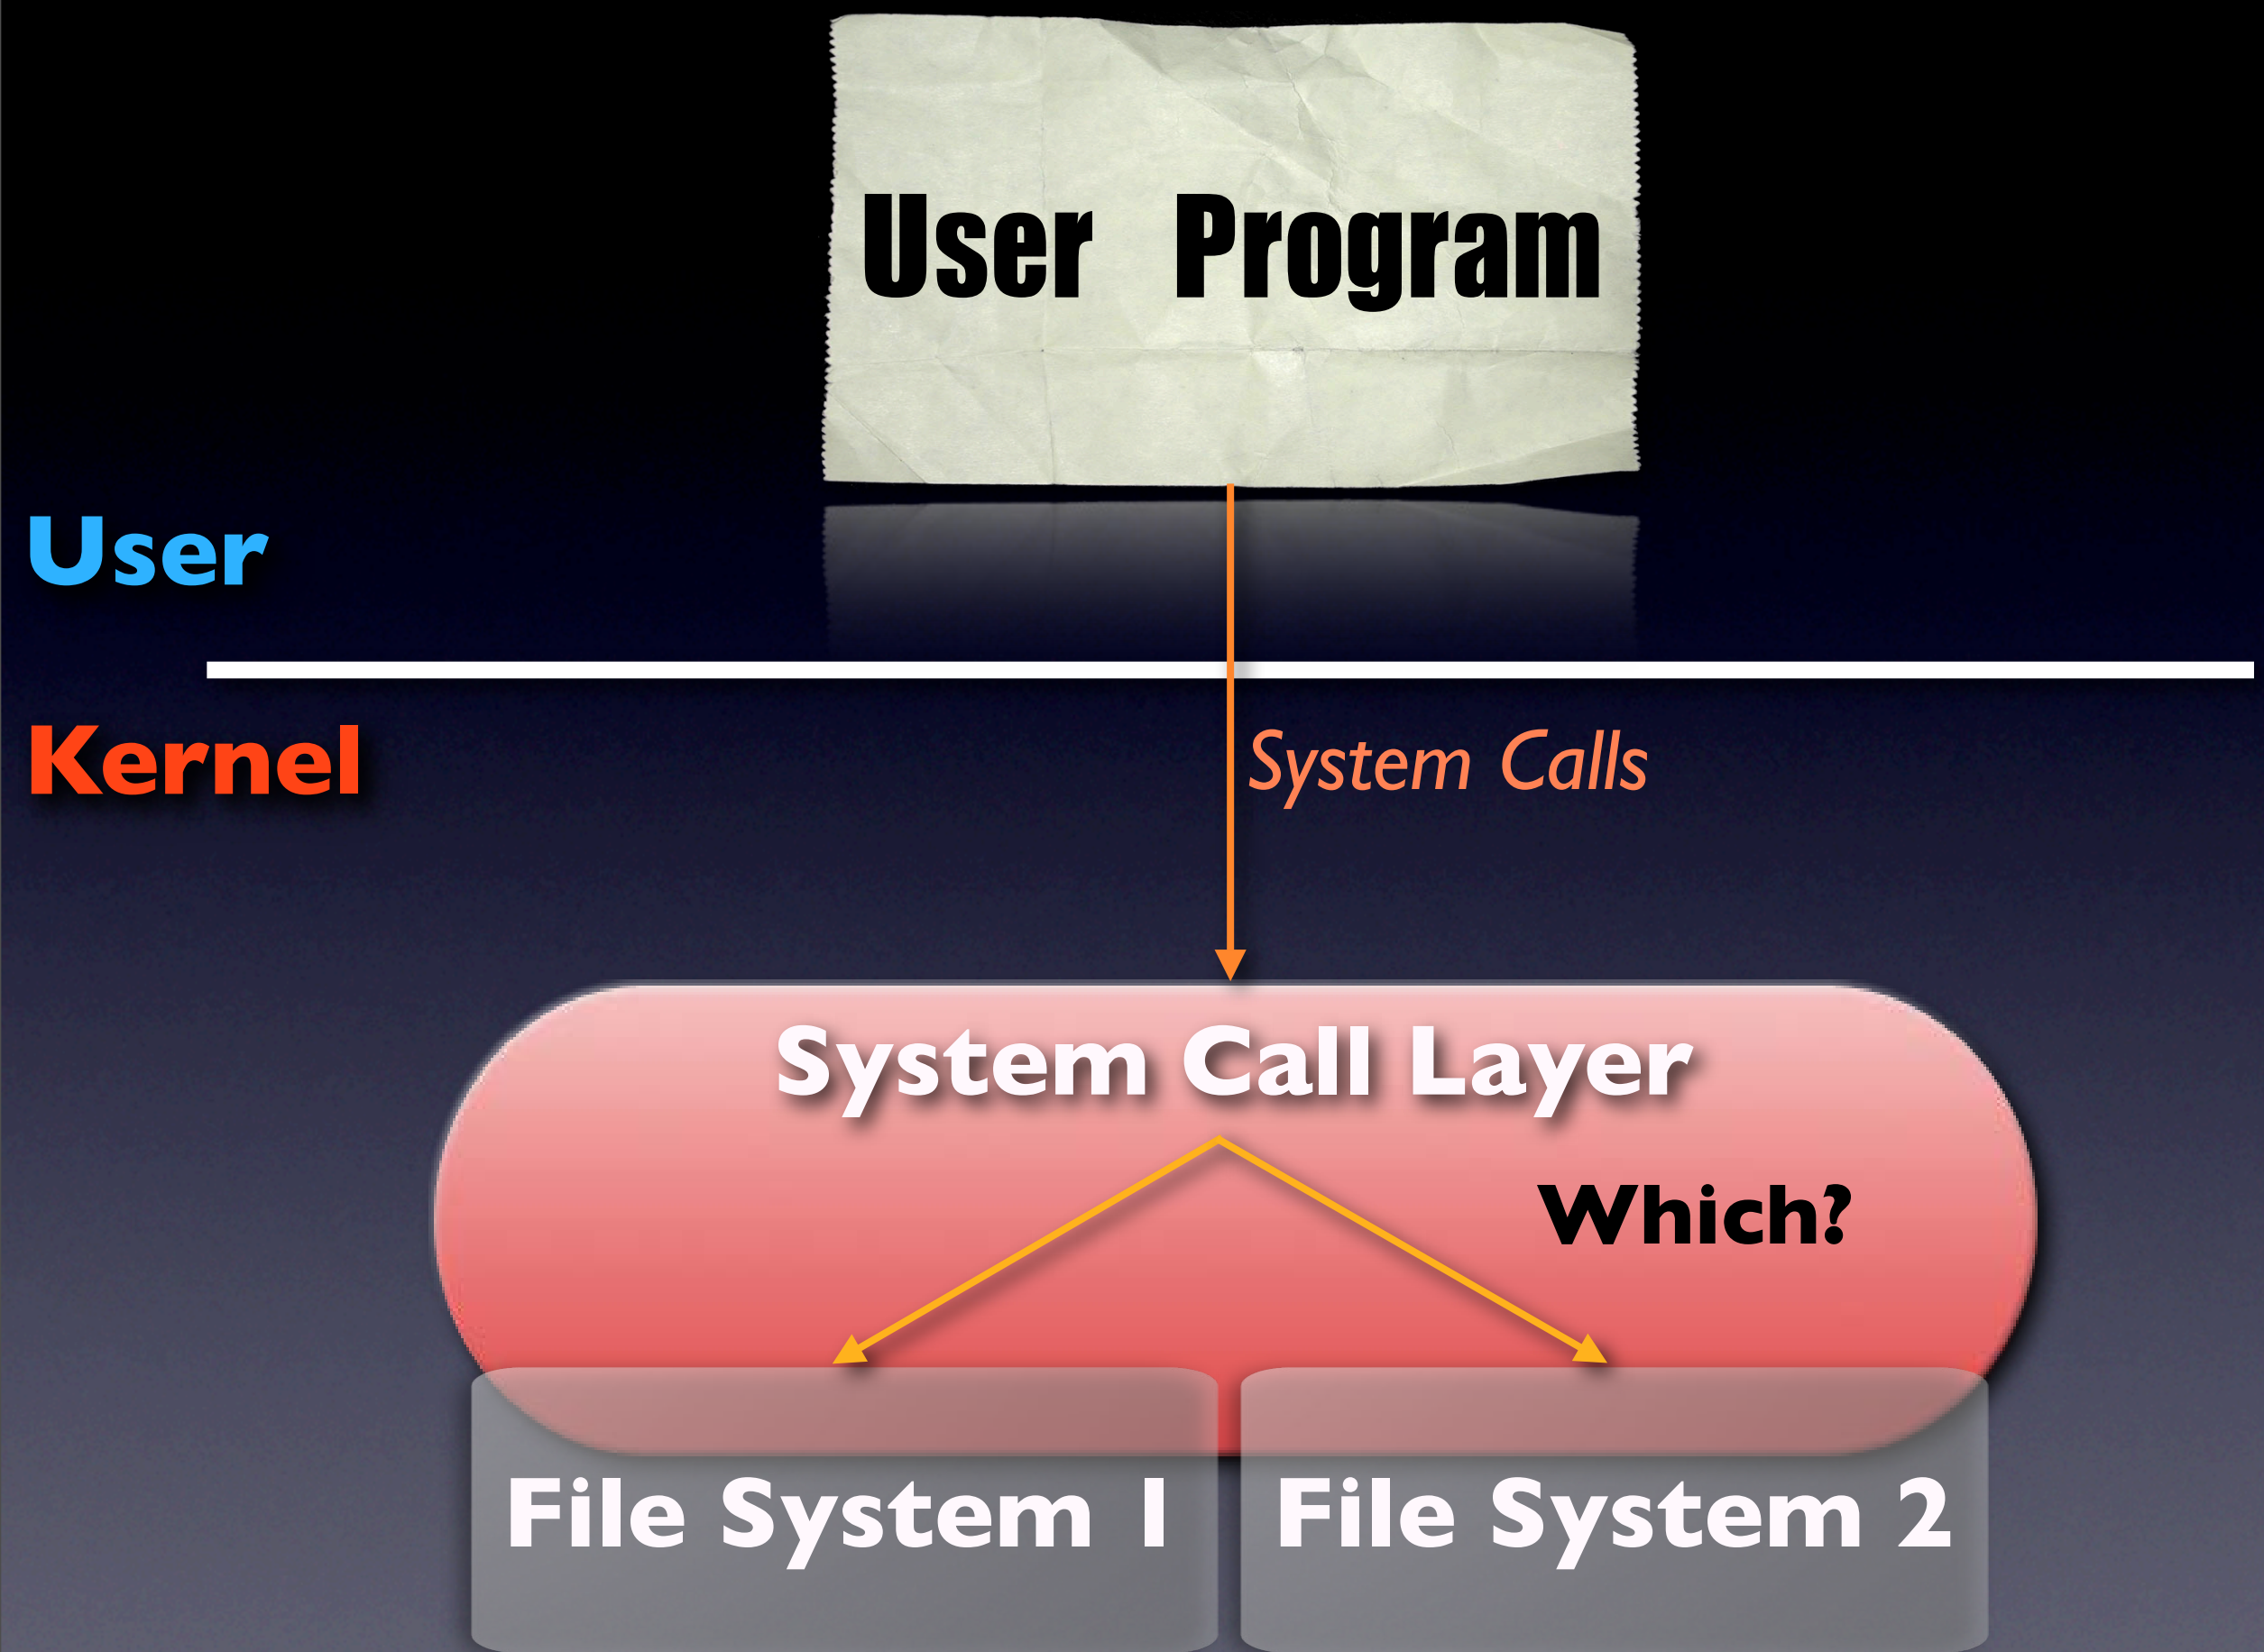
System Calls

System Call Layer

Which?

File System 1

File System 2



User Program

User

Kernel

System Calls

System Call Layer

Which?

FS₁

FS₂

FS₃

...

FS_n

User

Kernel

node₁

node_n

...

Create()

Remove()

Read()

Write()

...

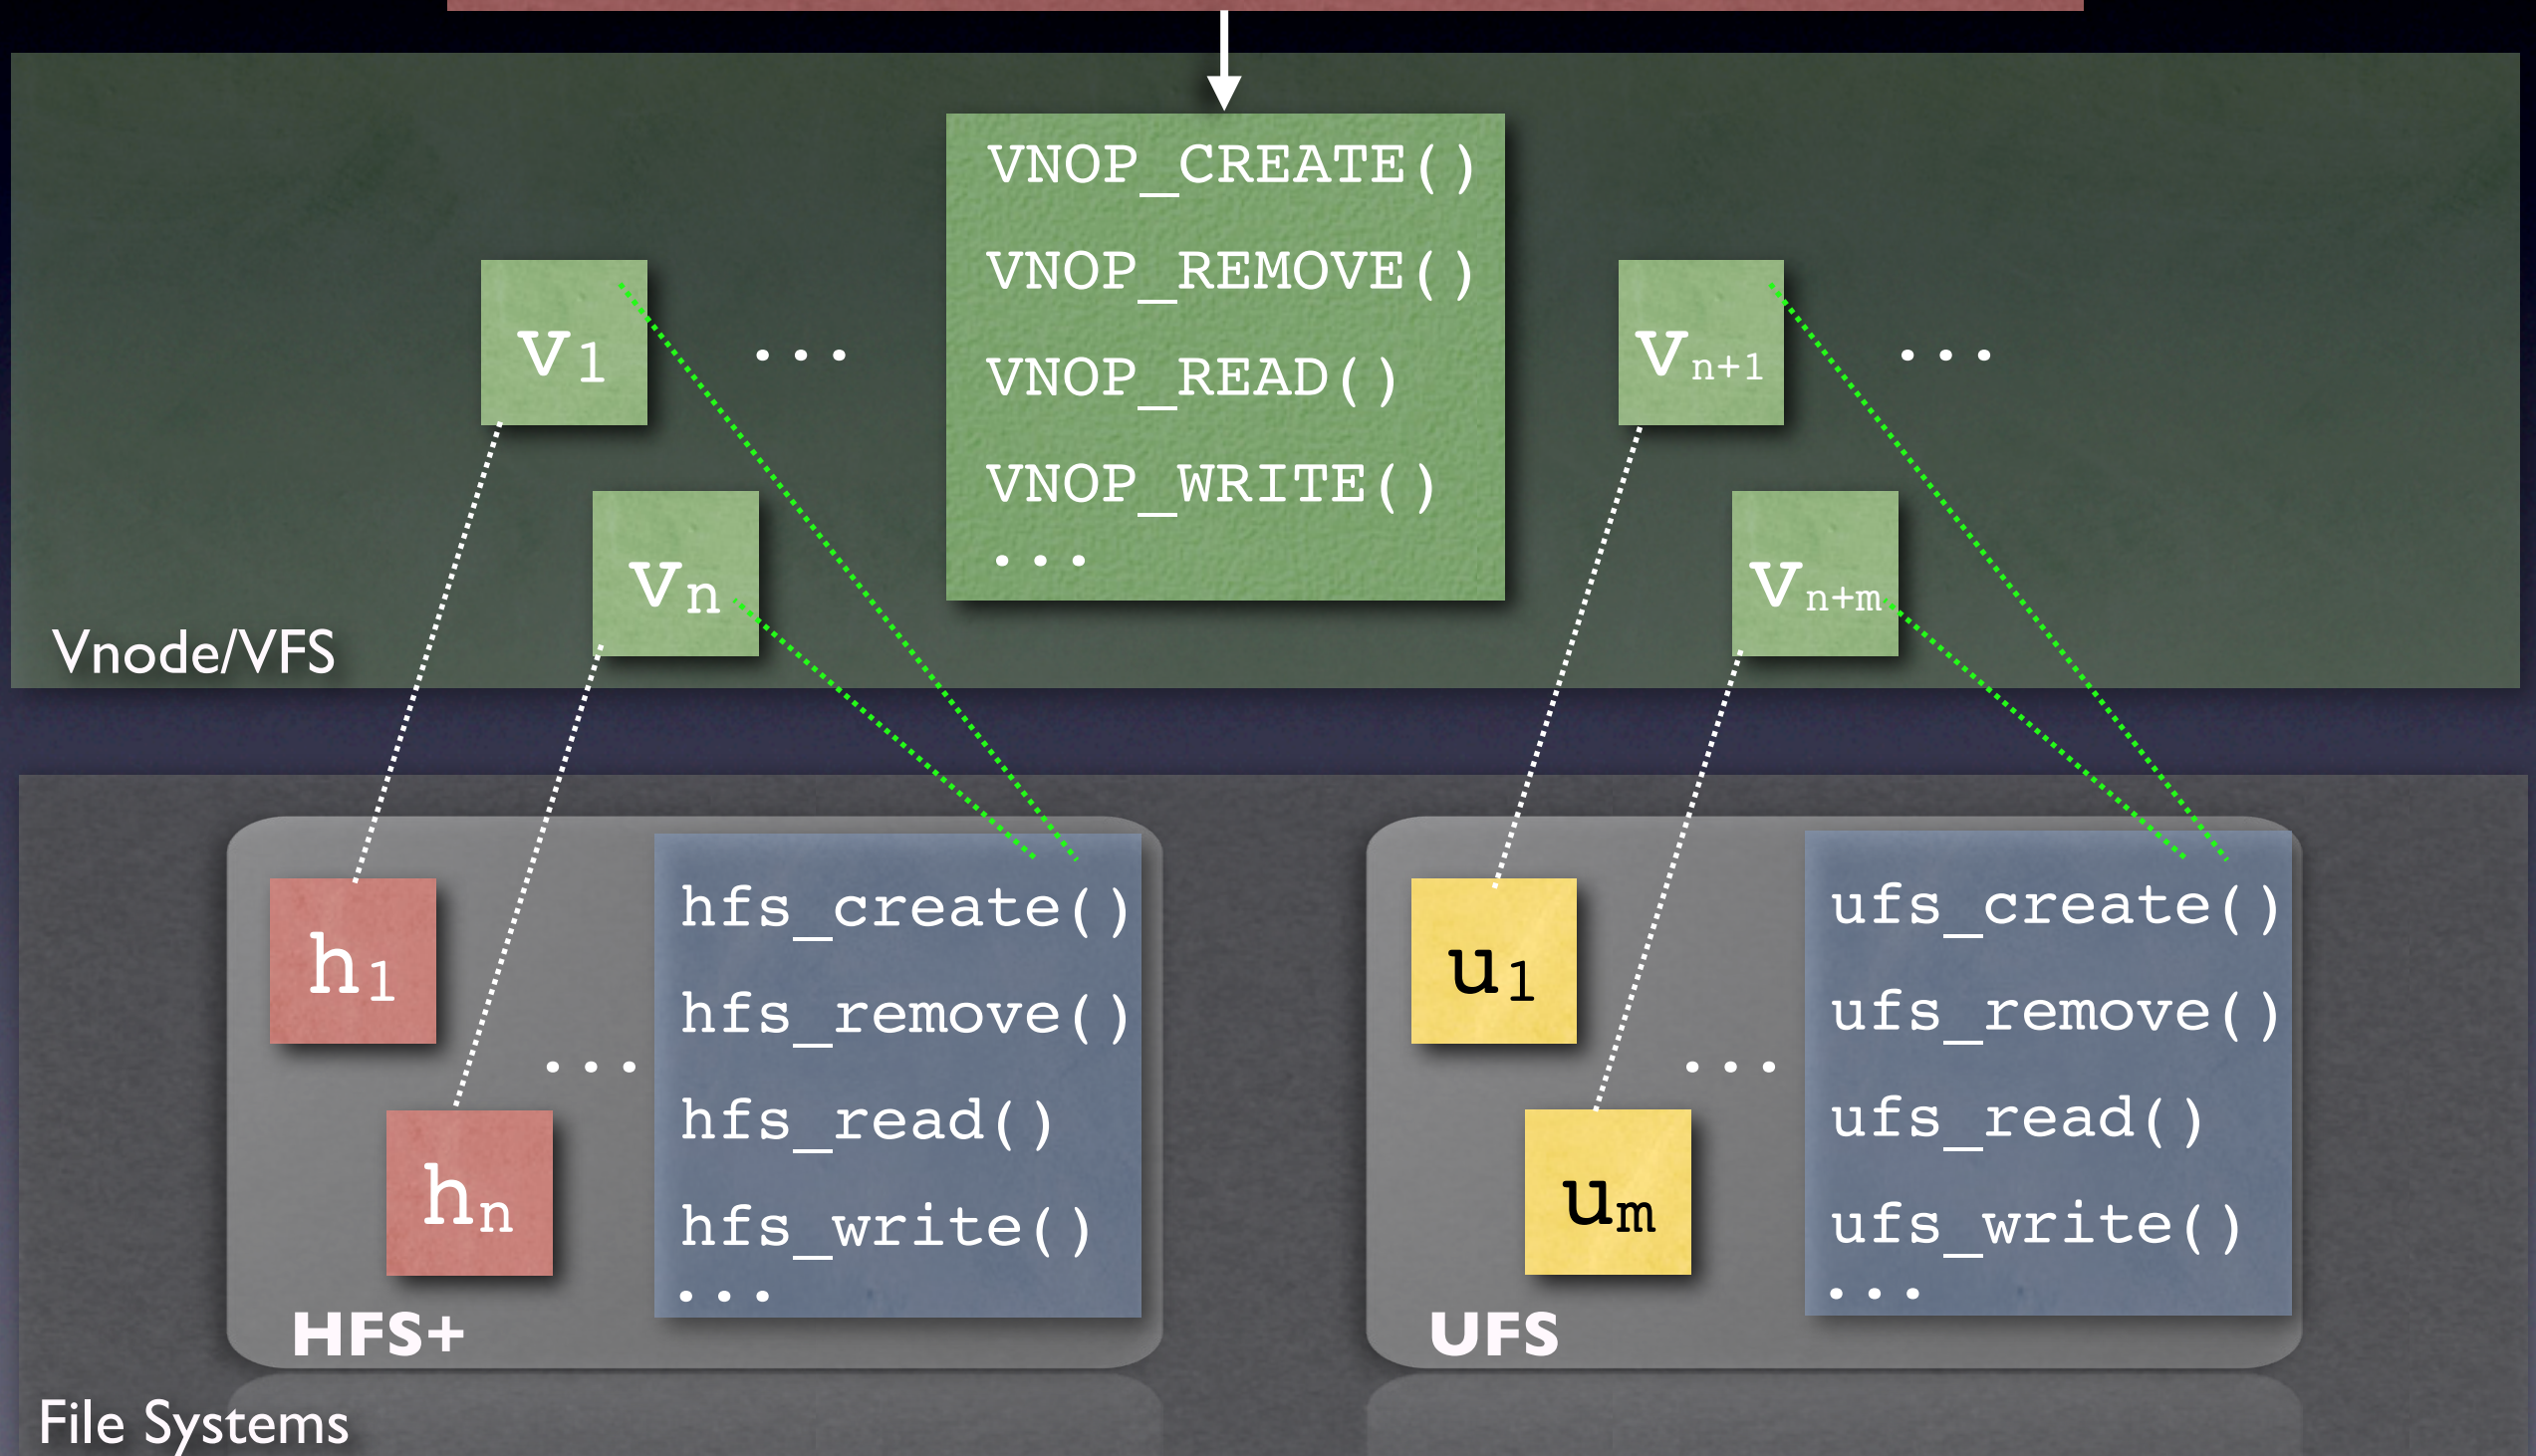
Some File System

User

User Program

Kernel

System Call Layer



HFS+

HFS

UFS

UDF

ISO9660

NTFS

MS-DOS

AFP

SMB/CIFS

NetFS

WebDAV

STFS

cddaafs

devfs

fatfs

FS

deadfs

synthfs

fifofs

union

specfs

volfs

Many File Systems

ZFS

Malloc/Free Abstraction

No partitions (just storage pool)


Automatically grows/shrinks

Pool's storage shared across file systems

Transactional (no journal, no fsck)

Copy-on-write, checksummed

Inside HFS+



```
$ sudo hfsdebug -f -t 10
```

```
...
```

```
# Top 10 Files with the Most Extents
```

```
1          632          183.10      2163846  
Macintosh HD:/.Spotlight-V100/.store.db
```

```
2          618          187.24      2163845  
Macintosh HD:/.Spotlight-V100/store.db
```

```
...
```

```
Out of 2252467 non-zero data forks total,  
2252128 (99.98 %) have no fragmentation.
```



```
$ sudo hfsdebug -0 | sort -n +1
```

```
# Free Contig Starting@ Ending@ Space
```

```
...
```

```
101813          0x2e1097b  0x2e2972f  397.71MB
```

```
127220          0x48f0c   0x67fff   496.95MB
```

```
238304          0x2dca9cd 0x2e04cac 930.88MB
```

```
...
```

```
$ sudo hfsdebug -m
```

```
...
```


```
next allocation search at block = 0x2c0f7d6
```

```
next unused catalog node ID    = 4042316
```

```
file system write count        = 138418087
```

```
...
```


Demos



<http://osxbook.com/software/hfsdebug/>

File System Activity

Carbon App

Cocoa App

X11 App

Cocoa
NSFileManager

**“Terminal”
Commands**

Carbon
Carbon File Manager

BSD Interfaces (libSystem)

Vnode/VFS Layer

HFS+

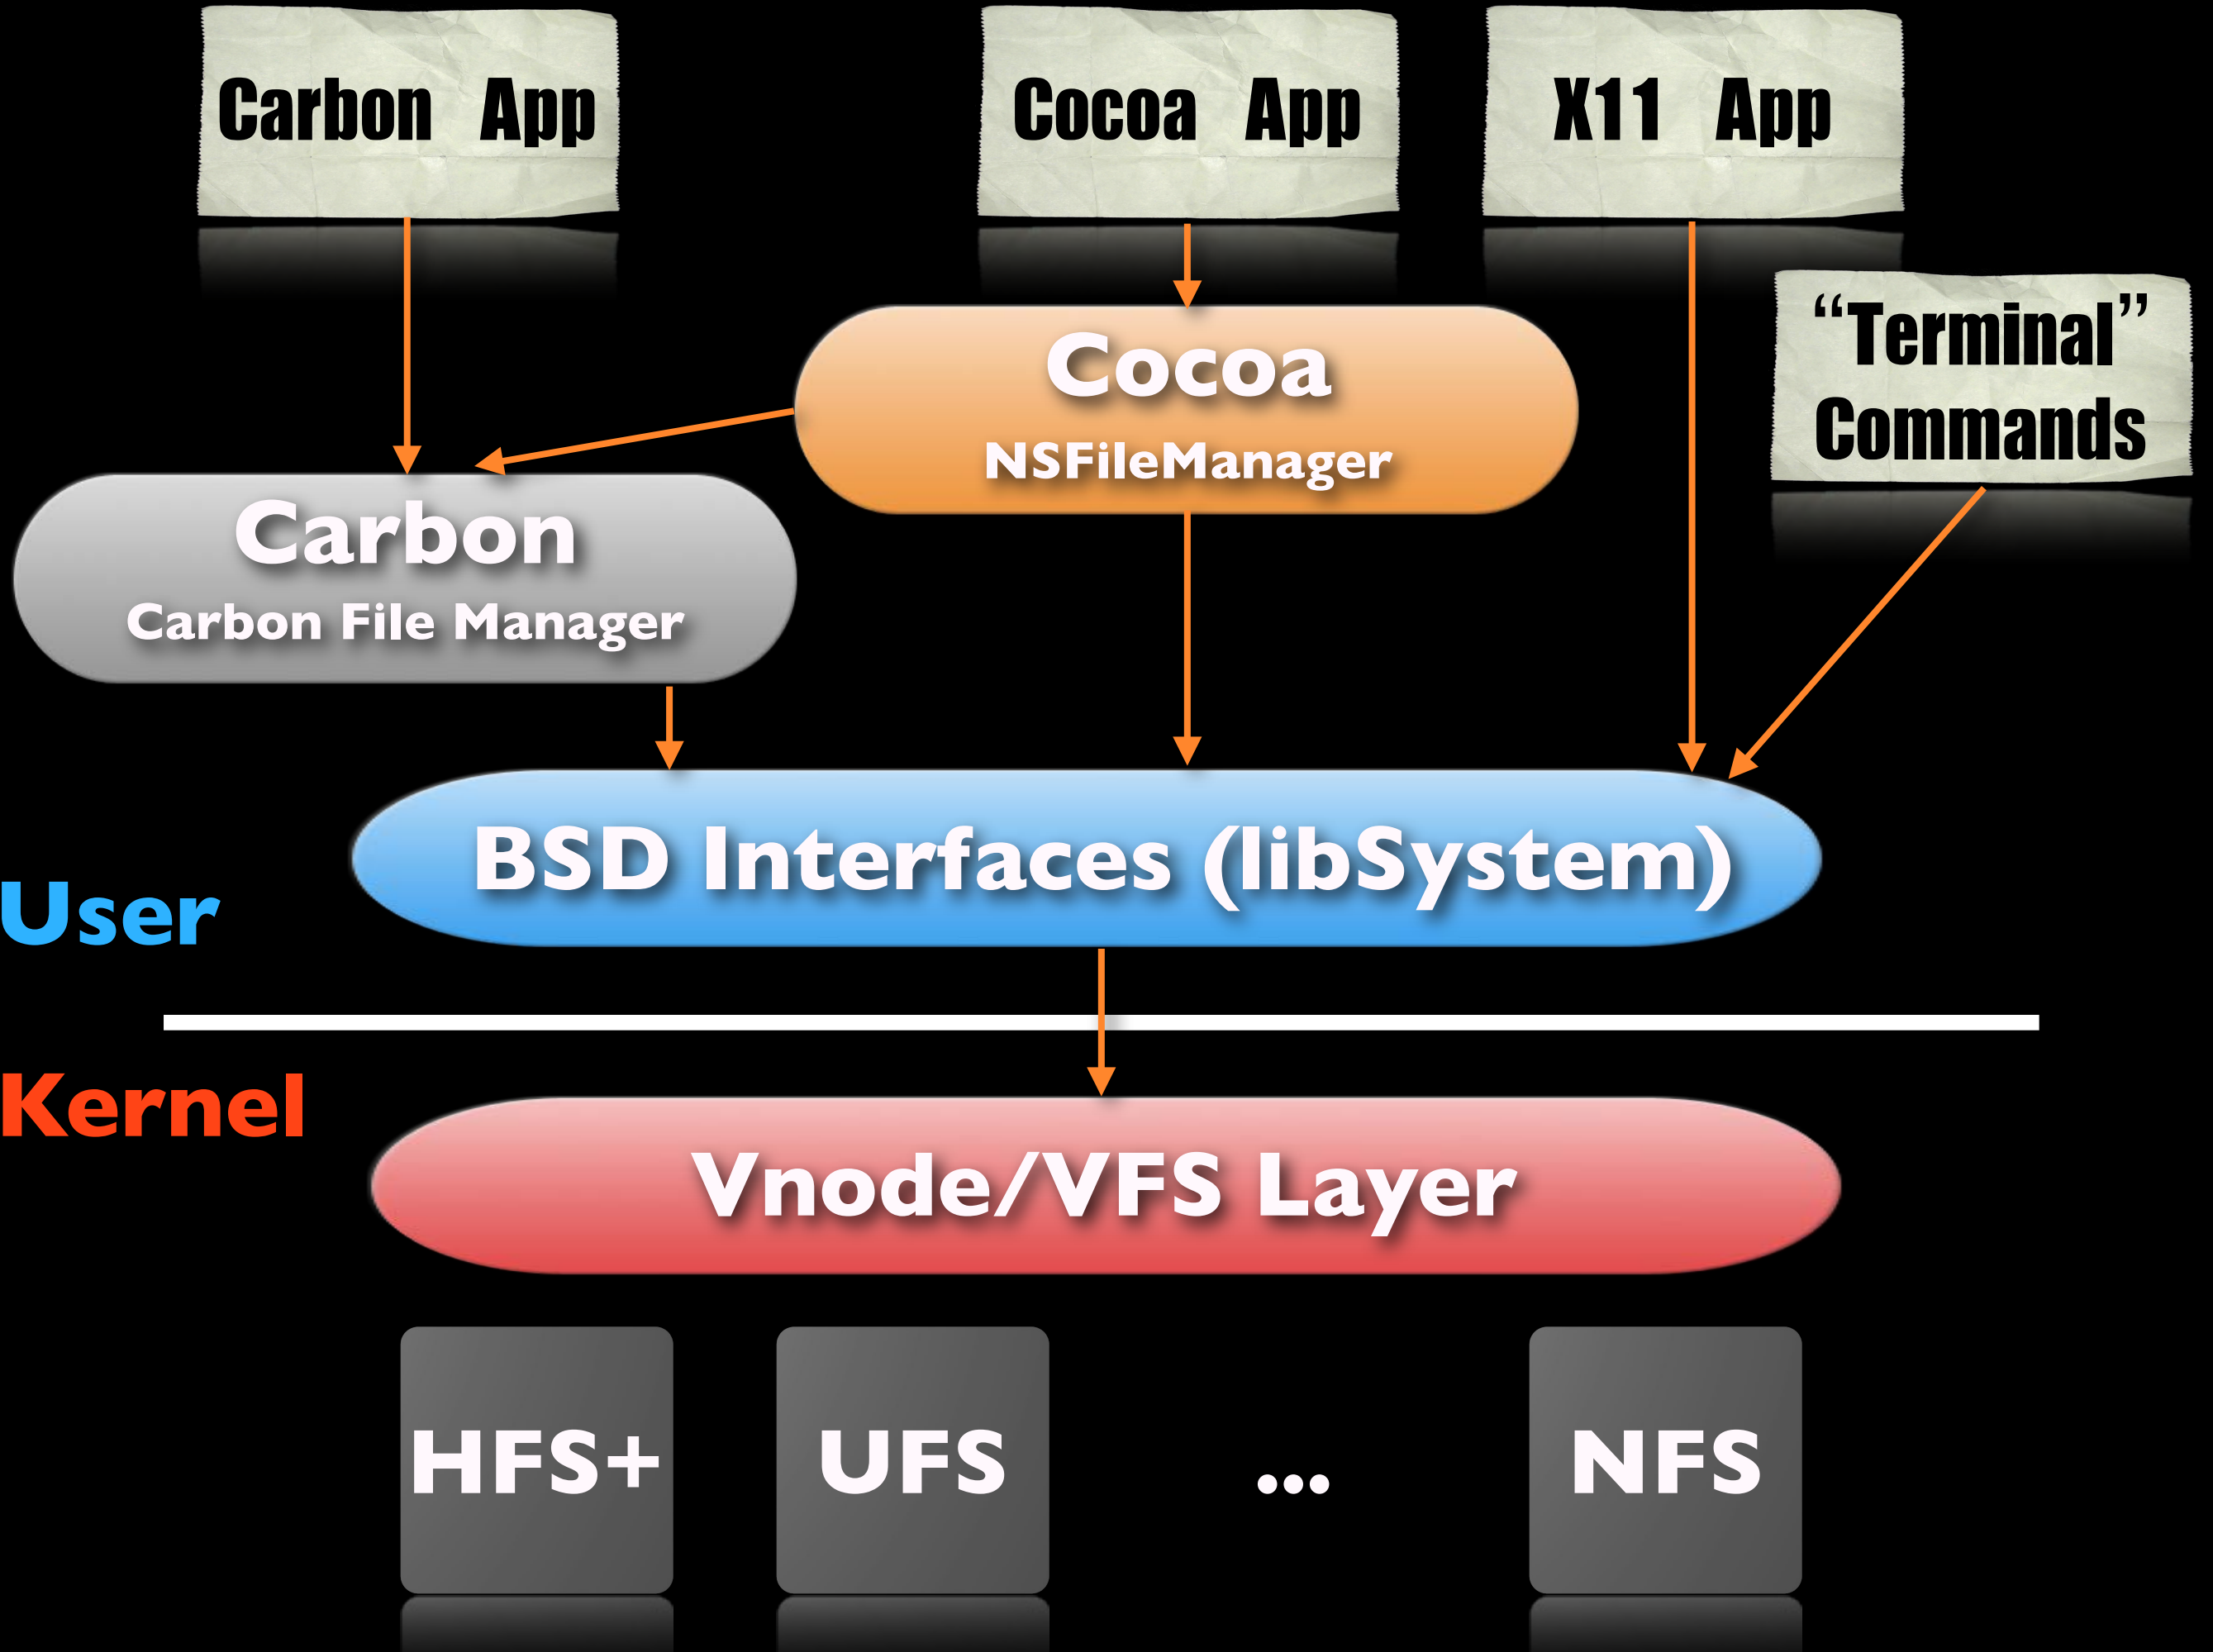
UFS

...

NFS

User

Kernel



Monitoring

File System Activity

fsevents

Kauth

kdebug

DTrace (Leopard onwards)



User Program

User

Kernel

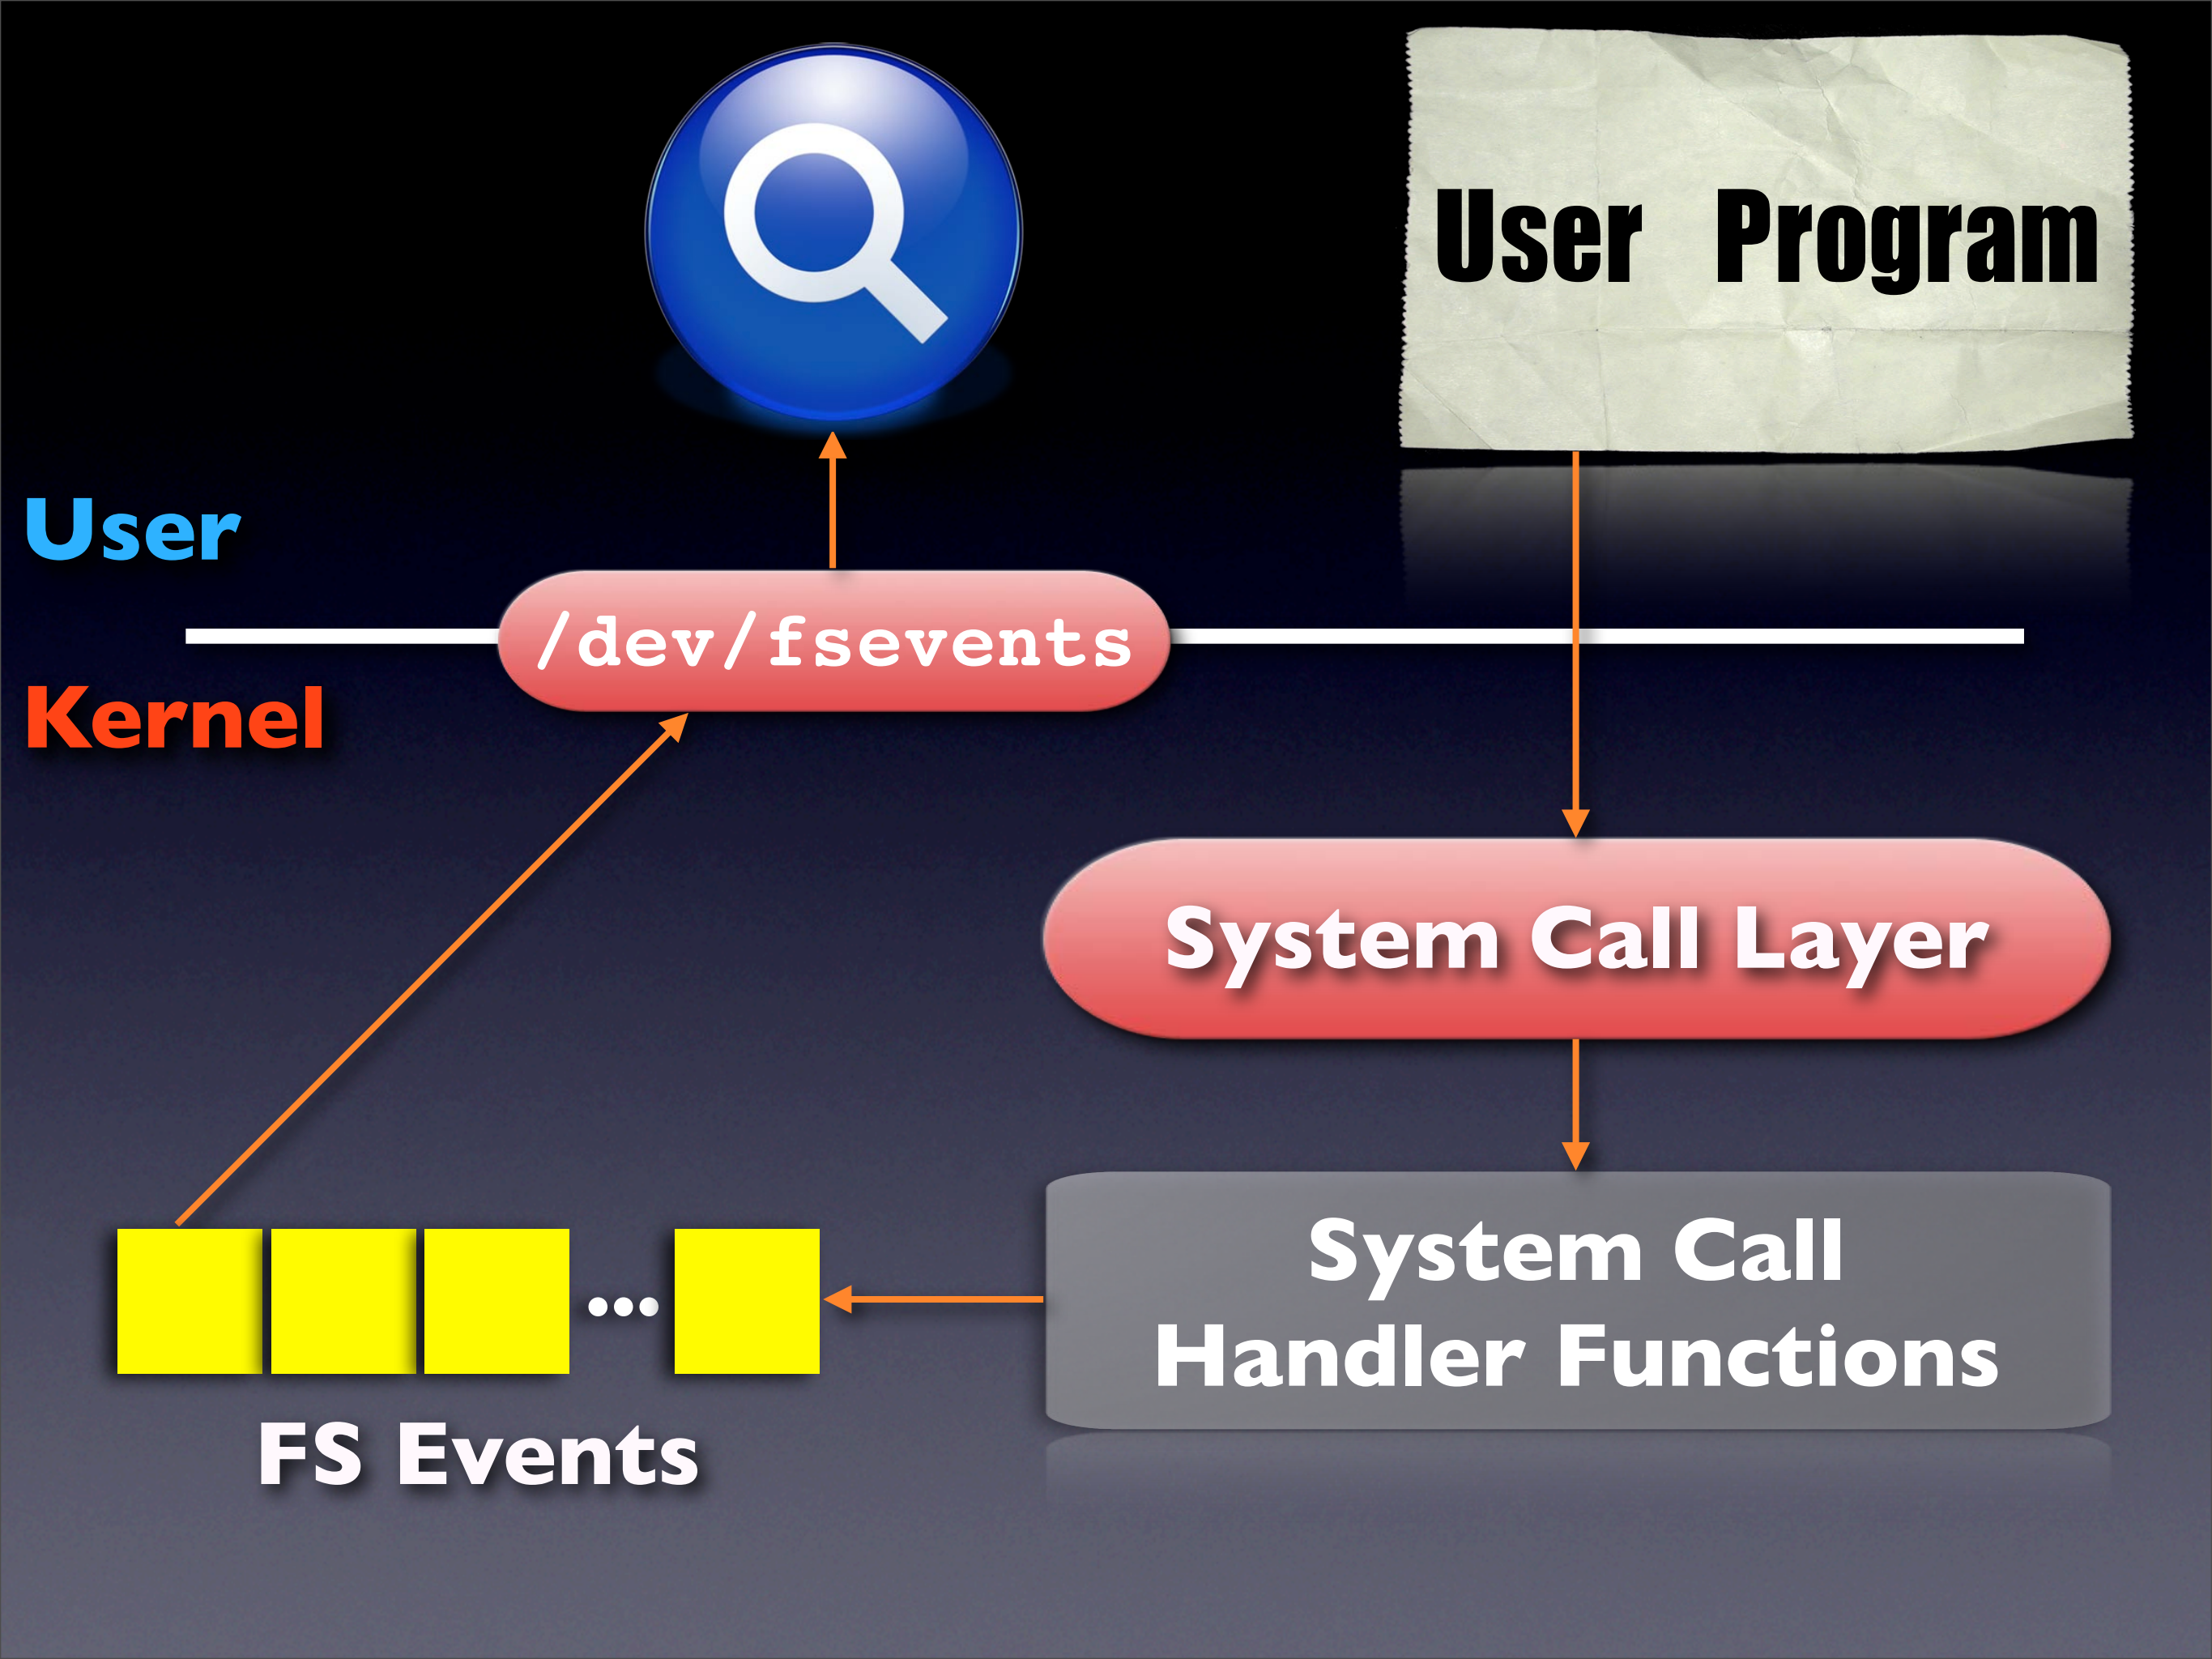
/dev/fsevents

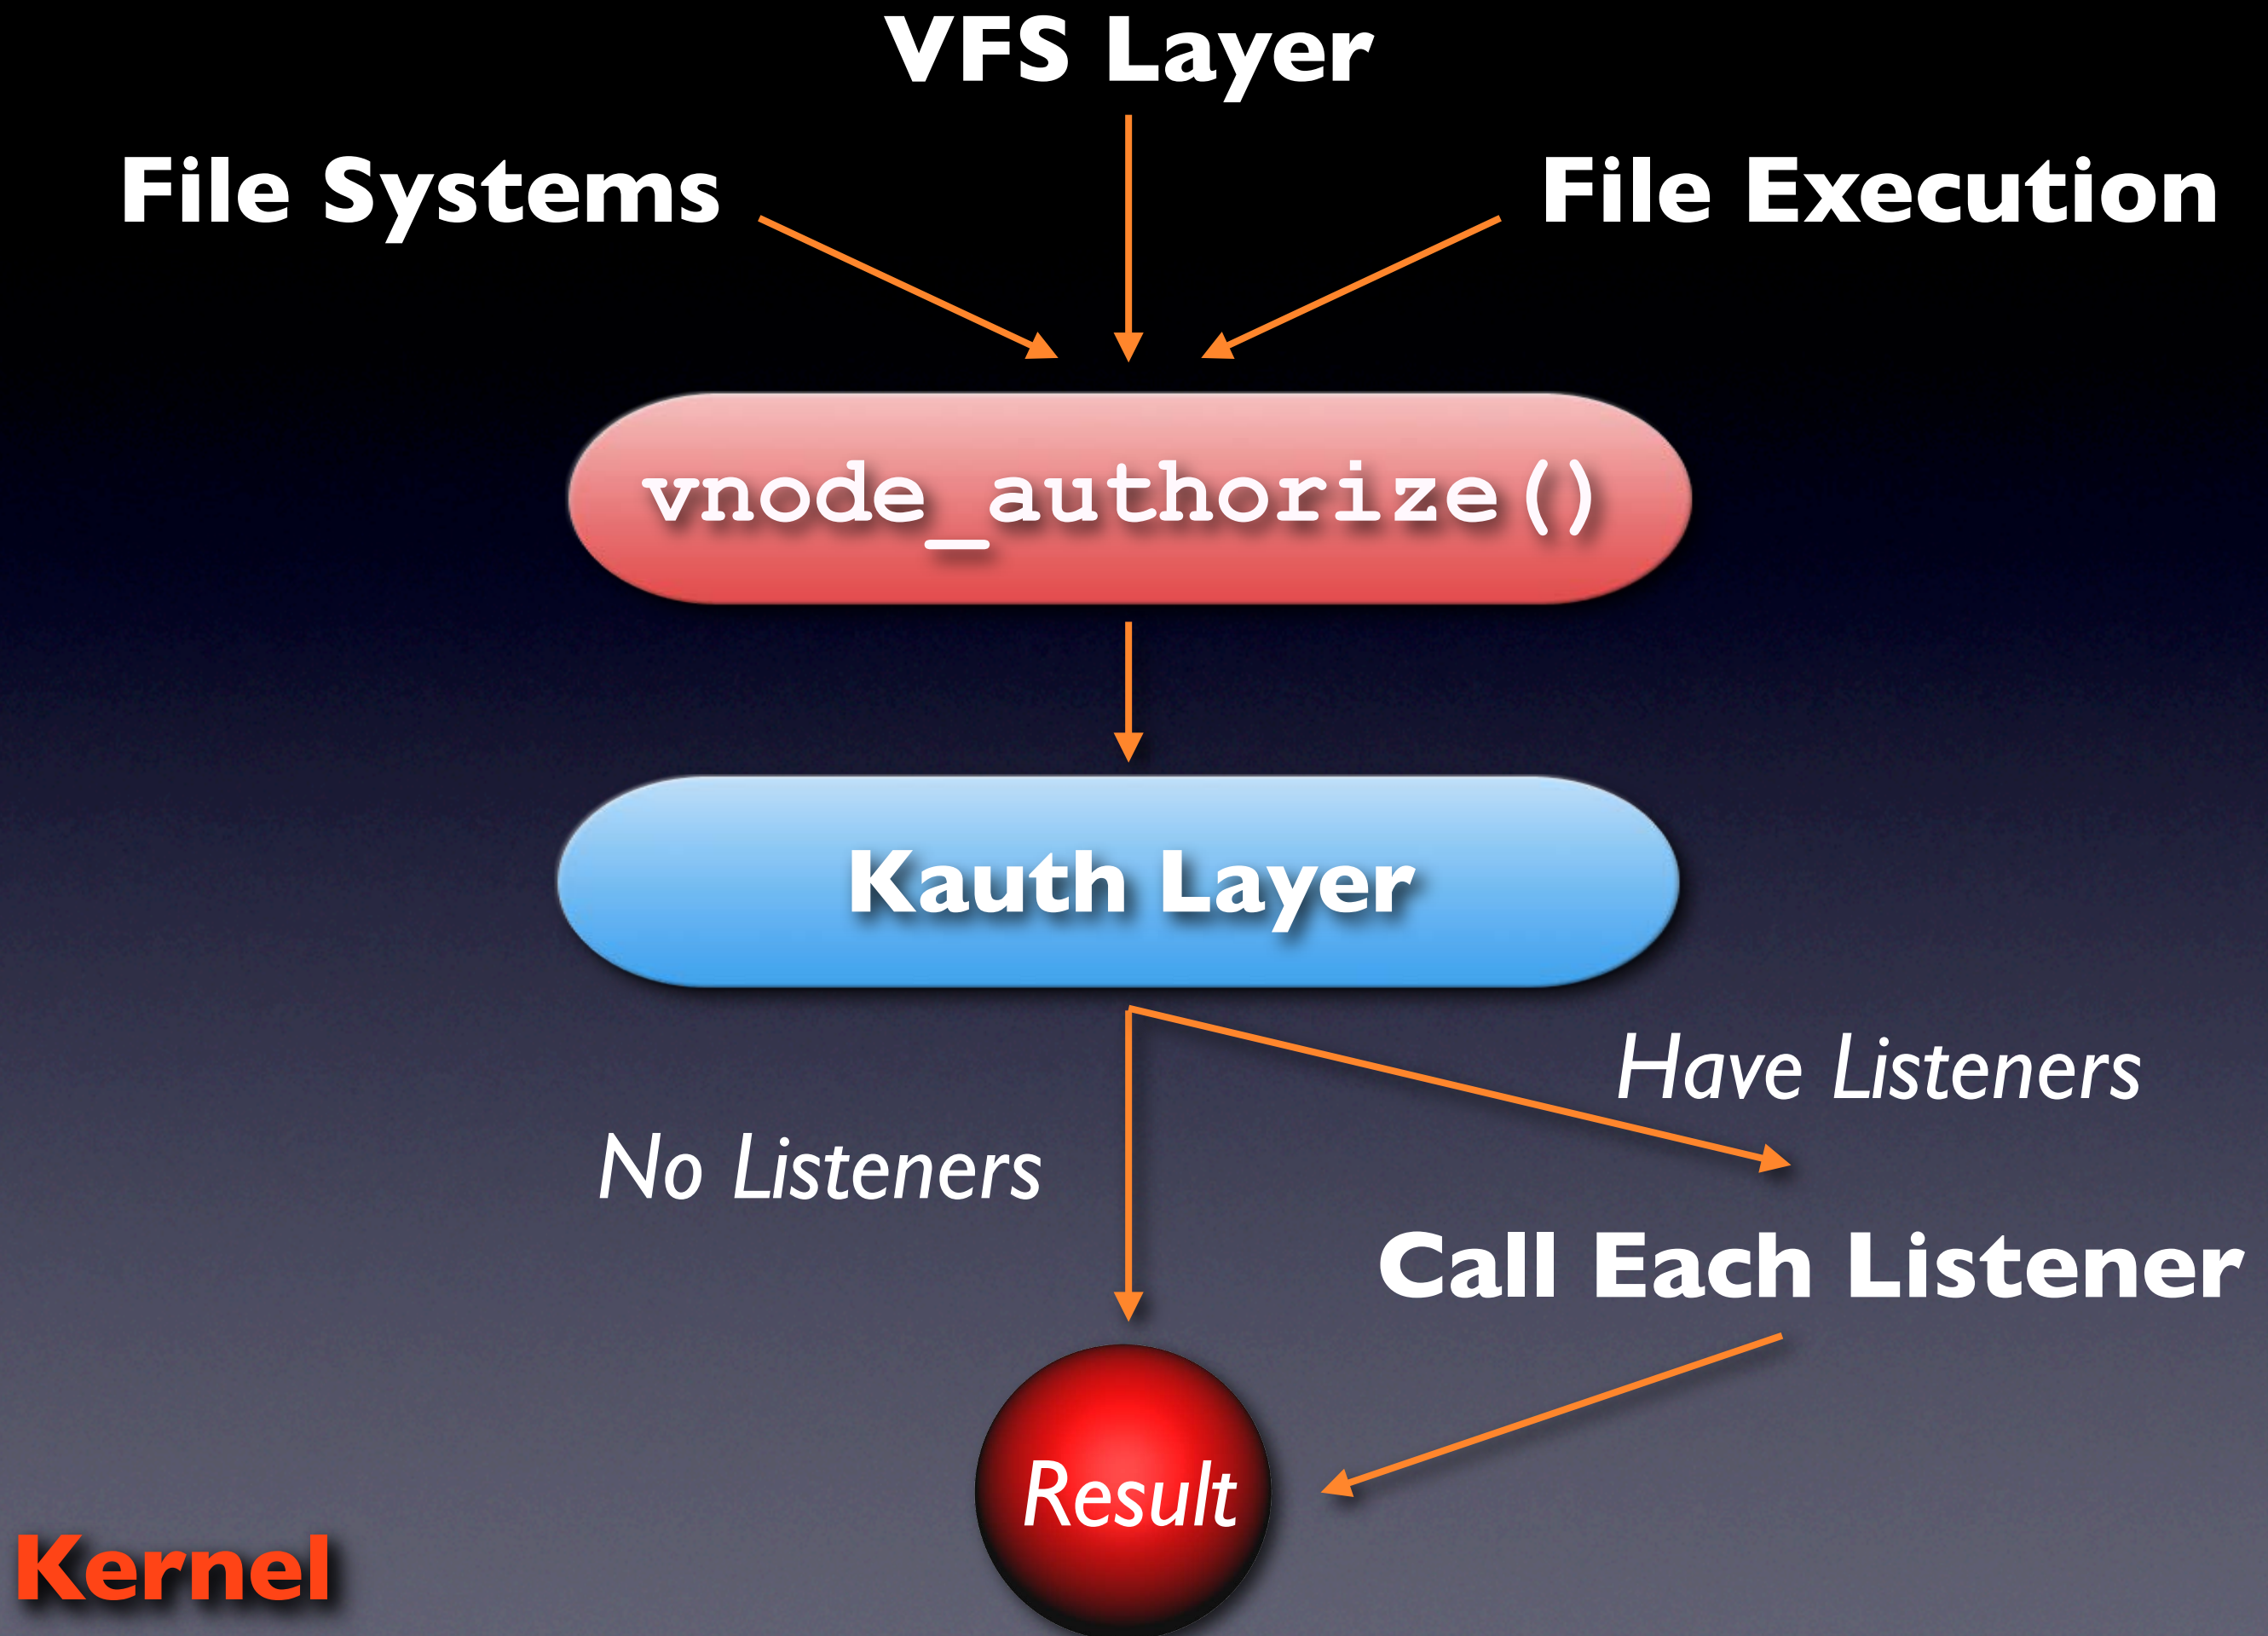
System Call Layer

**System Call
Handler Functions**



FS Events






fs_usage

latency

sc_usage

User

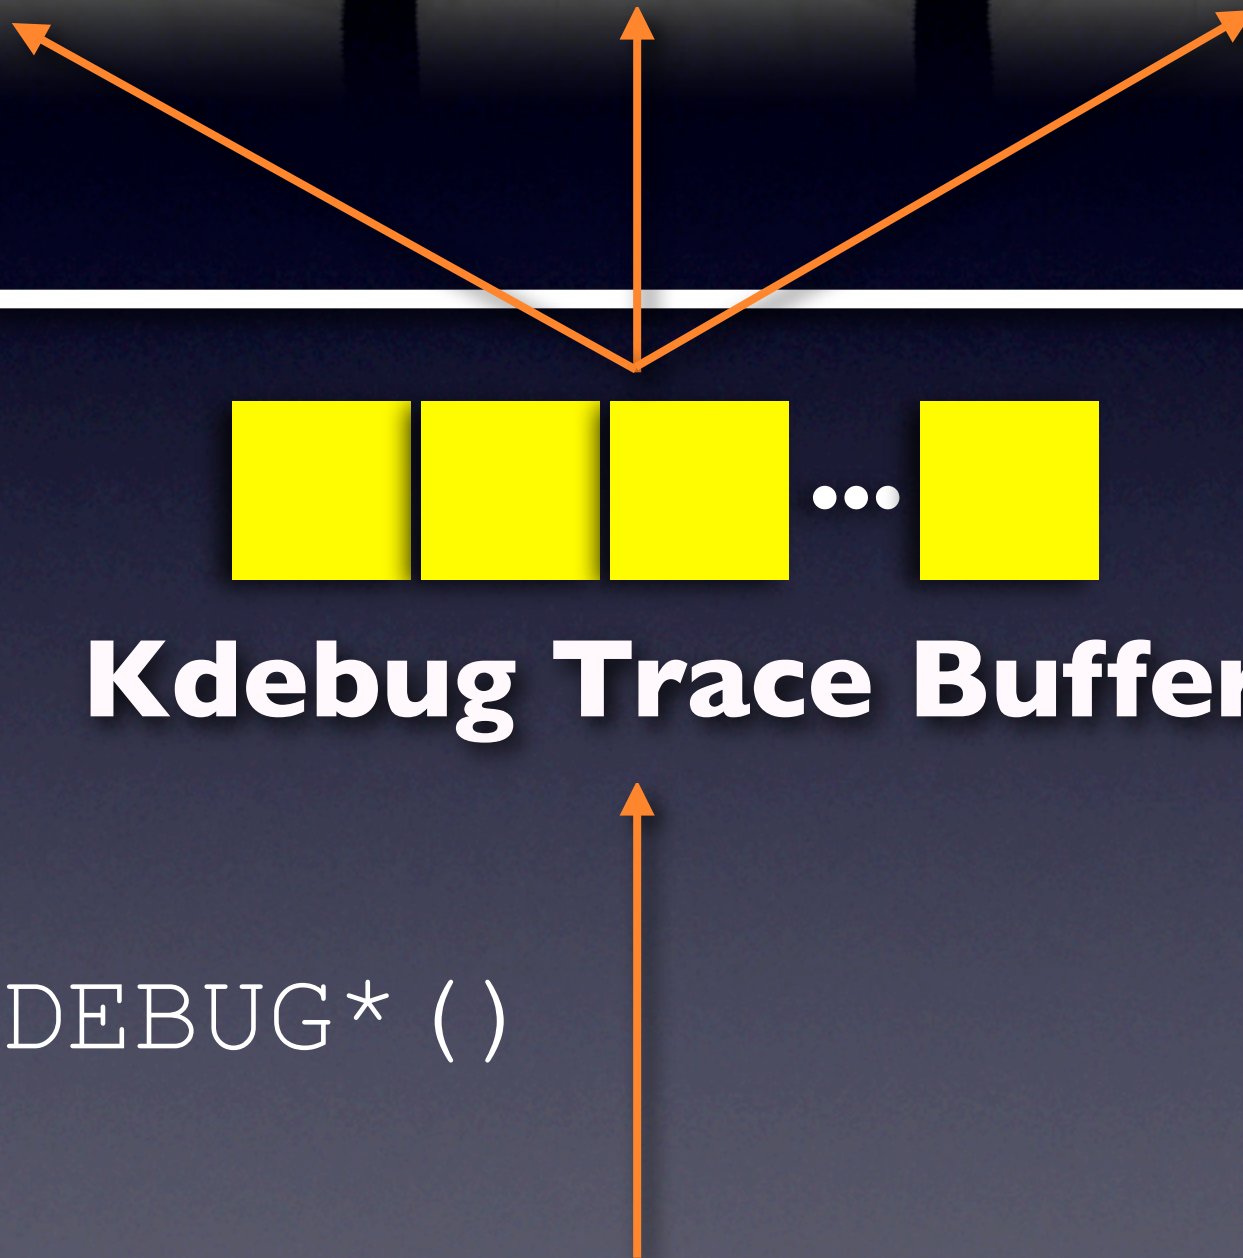
Kernel



Kdebug Trace Buffer

`KERNEL_DEBUG* ()`

Kernel Code



Modifying

File System Activity

Permissions and ACLs

Kauth

Storage Filter Scheme

User Program

Write

Read

...

Filter

Device Services

Transport Services


Physical Interconnect

**Your
Kernel
Extension**

Device

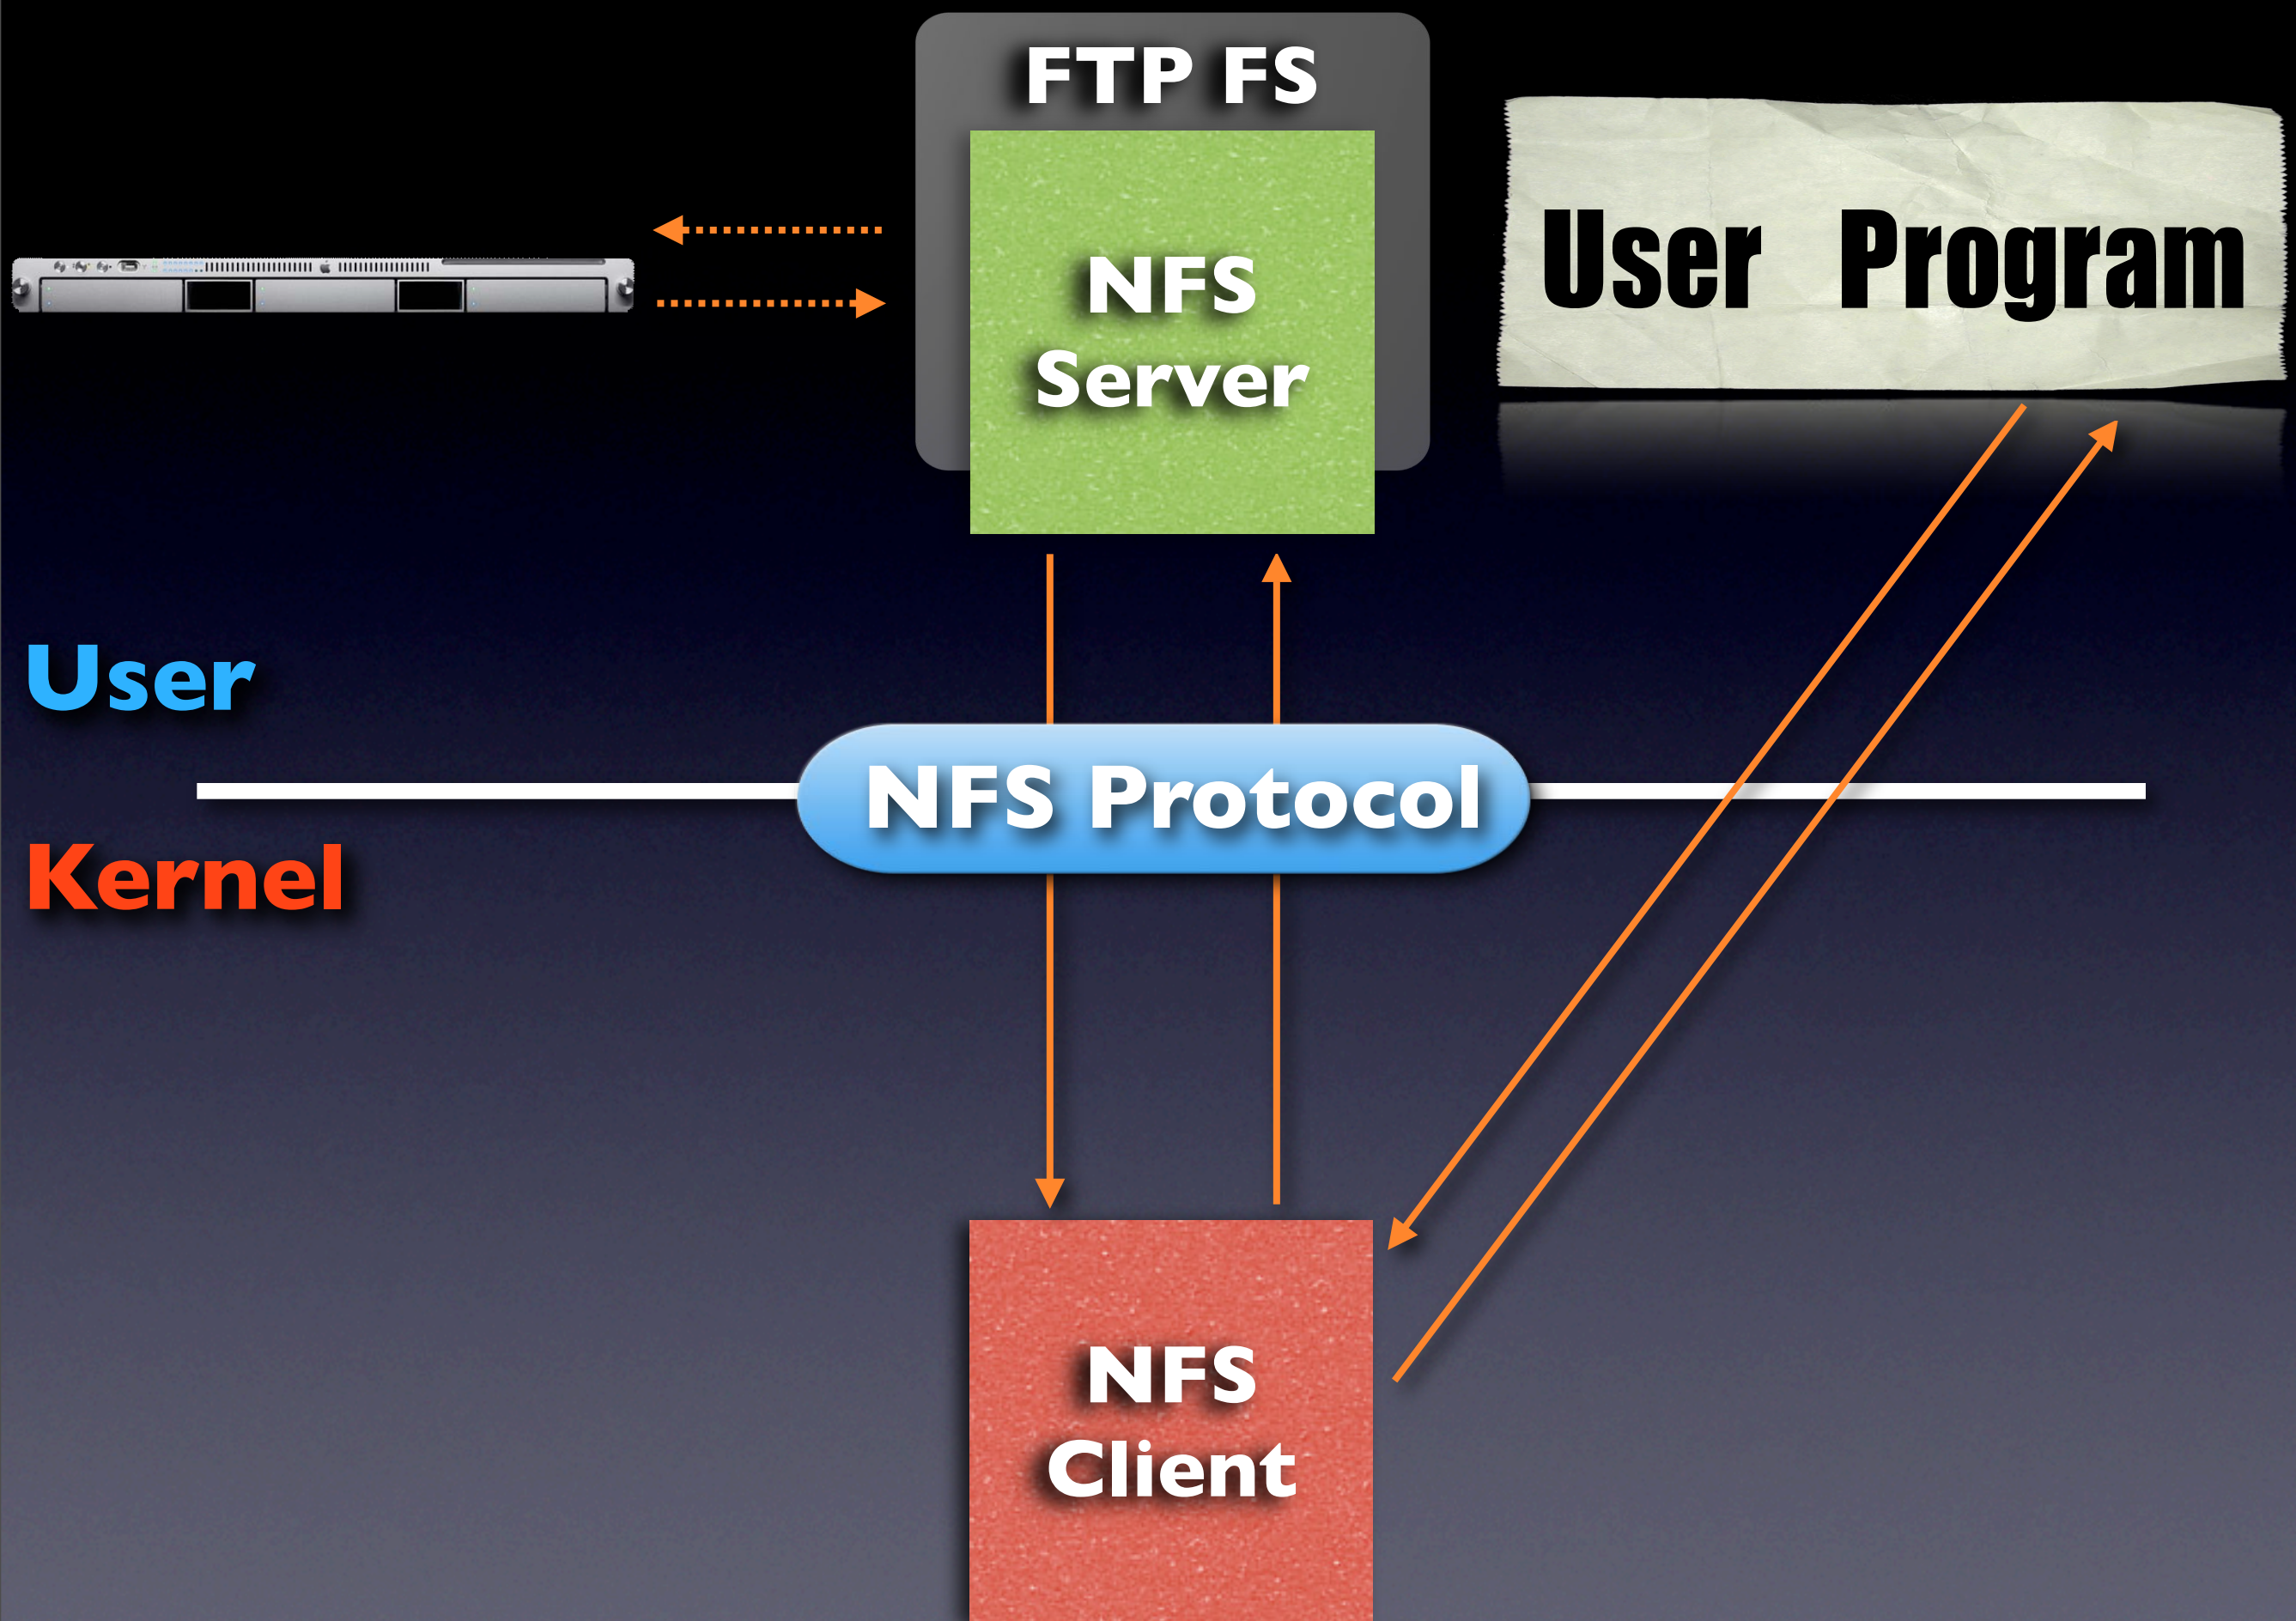


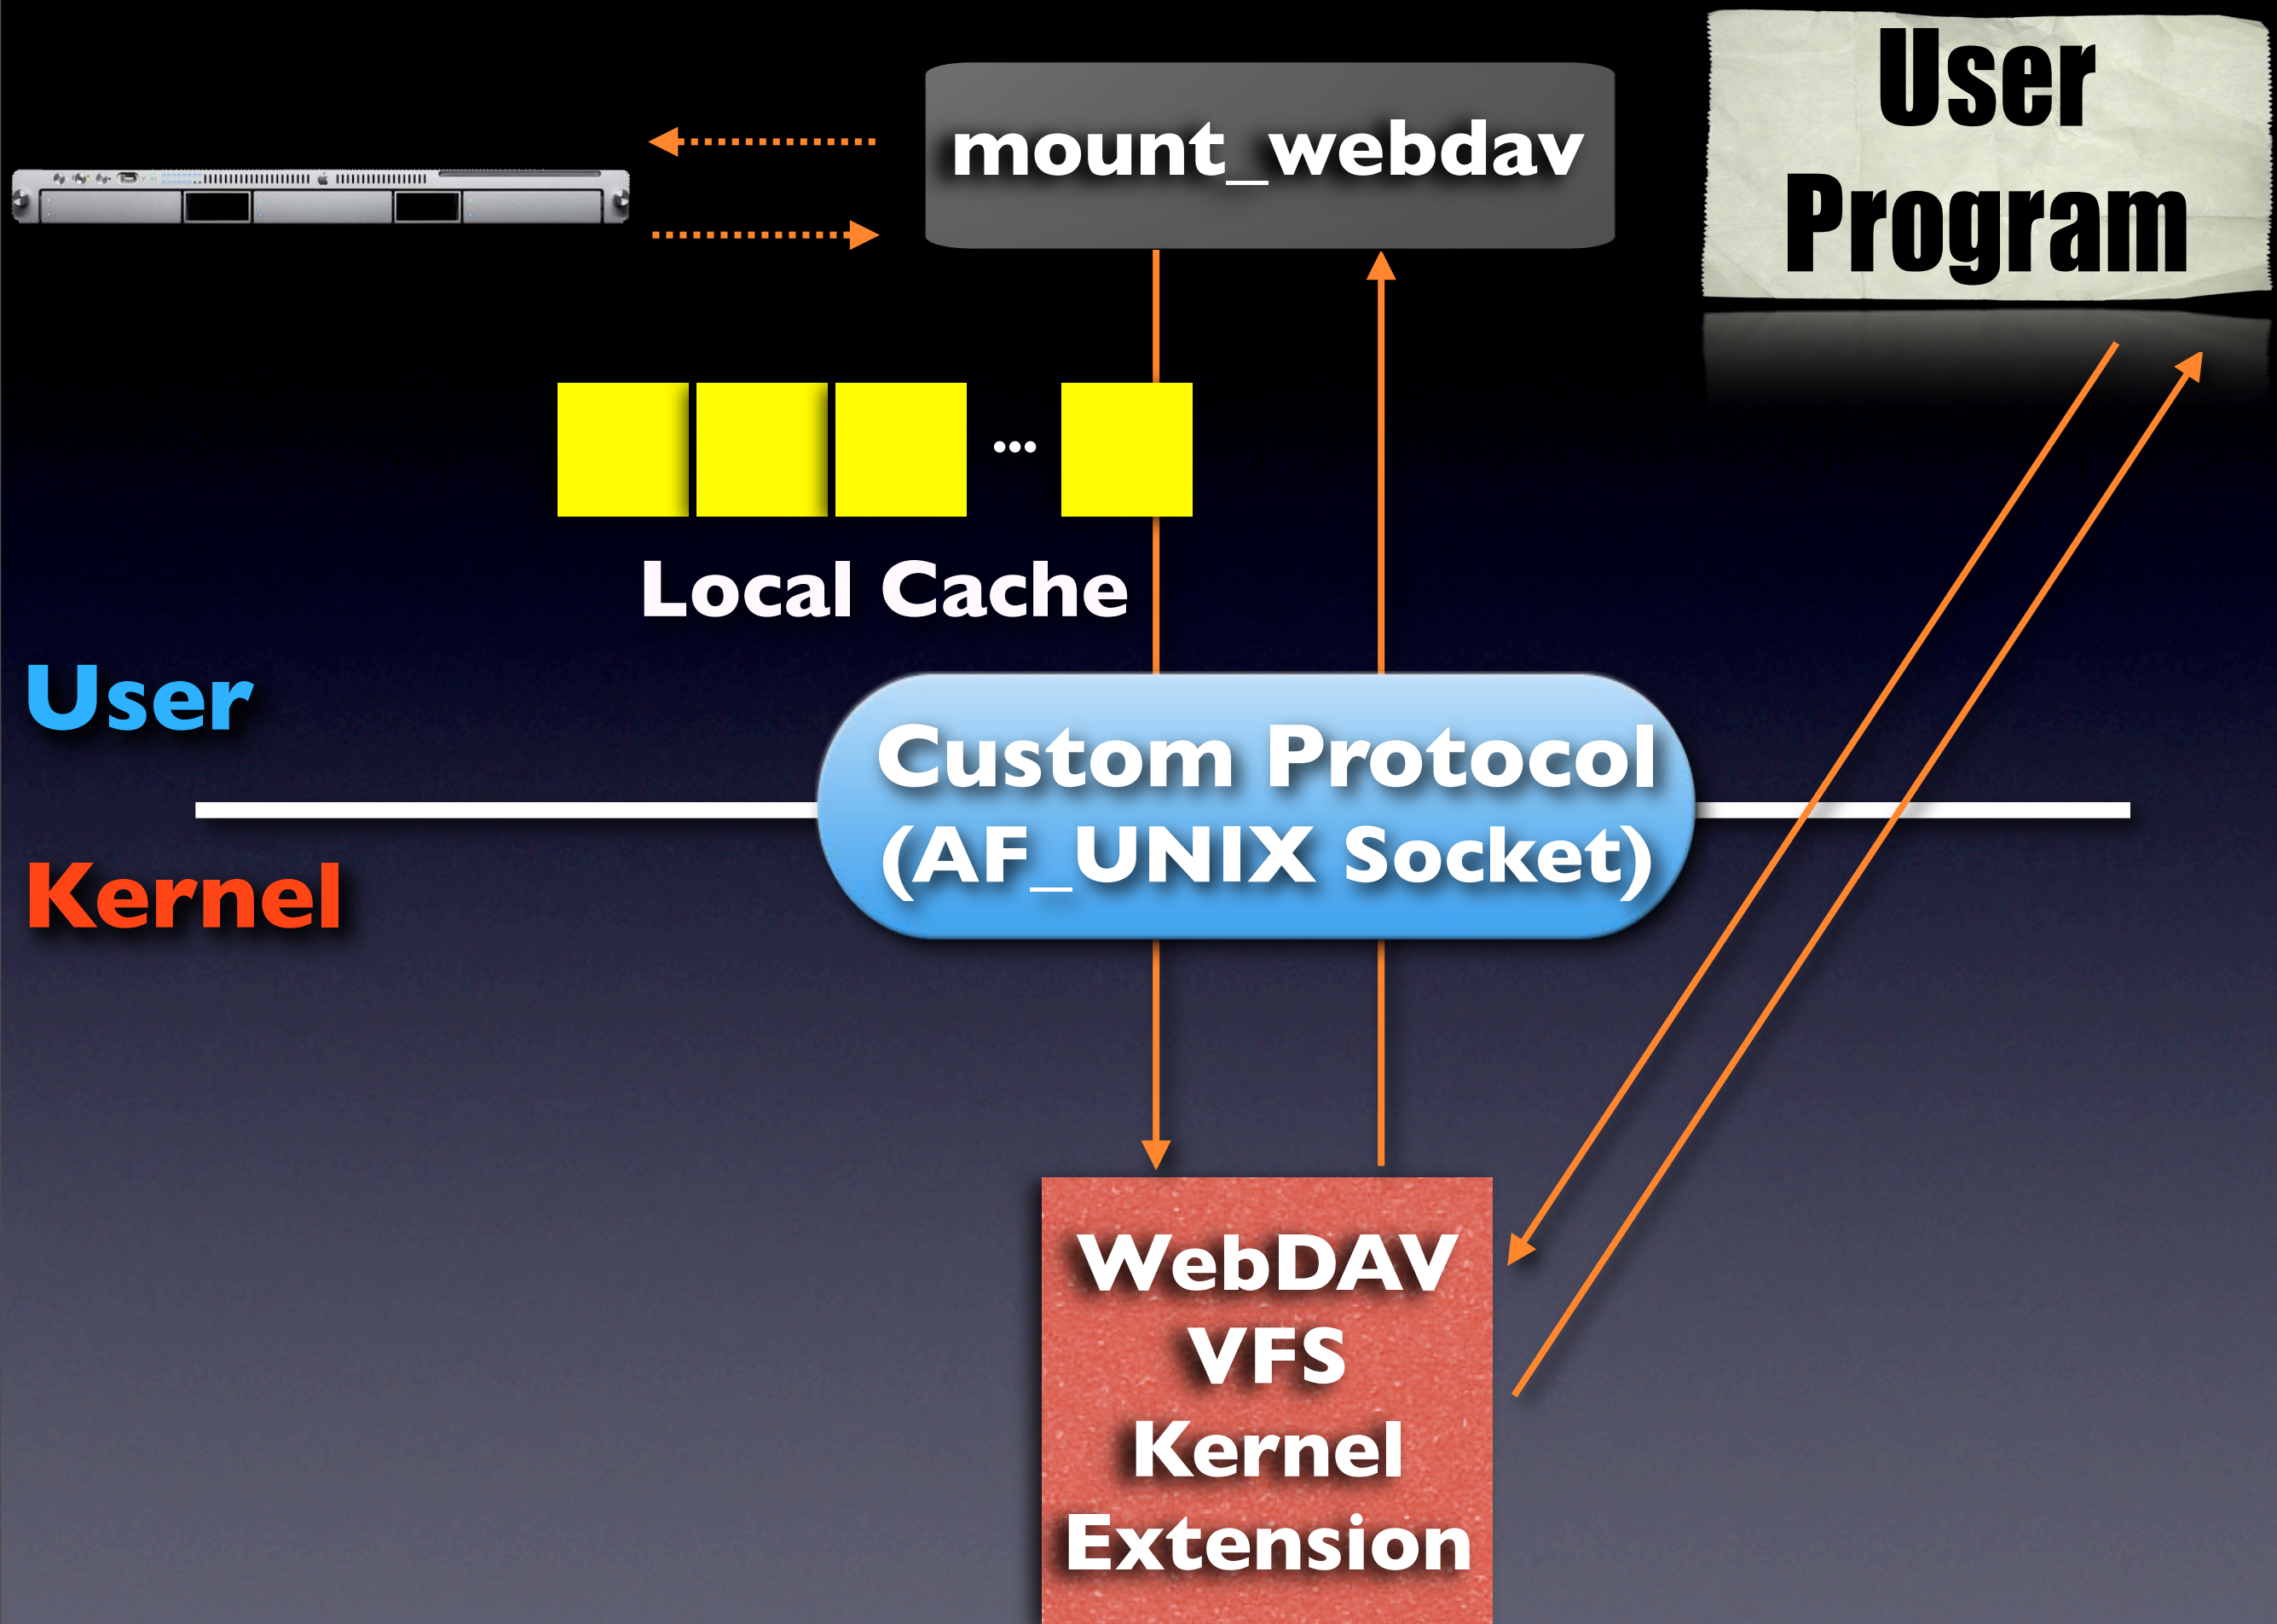
Demos



Hosting

File System Activity





Hosting

File System Activity

Difficult.

Until Today.

MacFUSE

**Source
of
Data**



**Your
File
System**

libfuse


User Program

User

Kernel


/dev/fuse

**MacFUSE
VFS**



FUSE&DPAP ferrisfuseFuseFTP unpackfs
beaglefs wdfs LAES gphoto2-fuse-fs
Bluetooth compFUSED gitfs spgmfs
fuseiso Logic FUR BlogFS FlickrFS
obexfs fsfin RelFS SMBNetFS ntfs-3g
ZFS sshfs GmailFS DBToy LZOLayer
ntfsmount FSFS PhoneBook GfarmFS NOOFS
hierfs Fusedav mcachefs Captive NTFS
owfs yacufs &ddfs Wikifuse fusedaap
unionfs-fuse Registered RAID TagsFS
fuse-dbfs KIO Fuse Gateway SMB for FUSE CopyFS
encfs Lkarmafs Cromfs FunionFS Mountlo
HTTP-FUSE rofs BitTorrent Itsvfs GSpaceFS
MythTVfs CvsFS Apertuiado GnomeVFS2FUSE
mediamax-fuse TrackerFS FunFS httpfs FUSEPod
Run-Time-Access djmount FUSE::DBI
gcfuse wafhs Hippocampus Wikipeidias
Sicfs User-Level Versioning Afuse
dereferencefs CryptoFS DVDfs PUFFS mysqlfs
LUFS Bridge fdmsfs archivemount
localfs CurlFtpFS mp3fs Offline Media Content DB
9pfuse

Demos



What's Being Released

MacFUSE VFS kernel extension

Mac OS X patch for libfuse

Mac OS X patch for sshfs

A Universal Binary tarball

<http://code.google.com/p/macfuse>

Thanks

FUSE Developers on Linux and FreeBSD

Google

My Team

You

Q/A

