



**TECHNISCHE  
UNIVERSITÄT  
DRESDEN**

Department of Computer Science Institute of System Architecture, Operating Systems Group

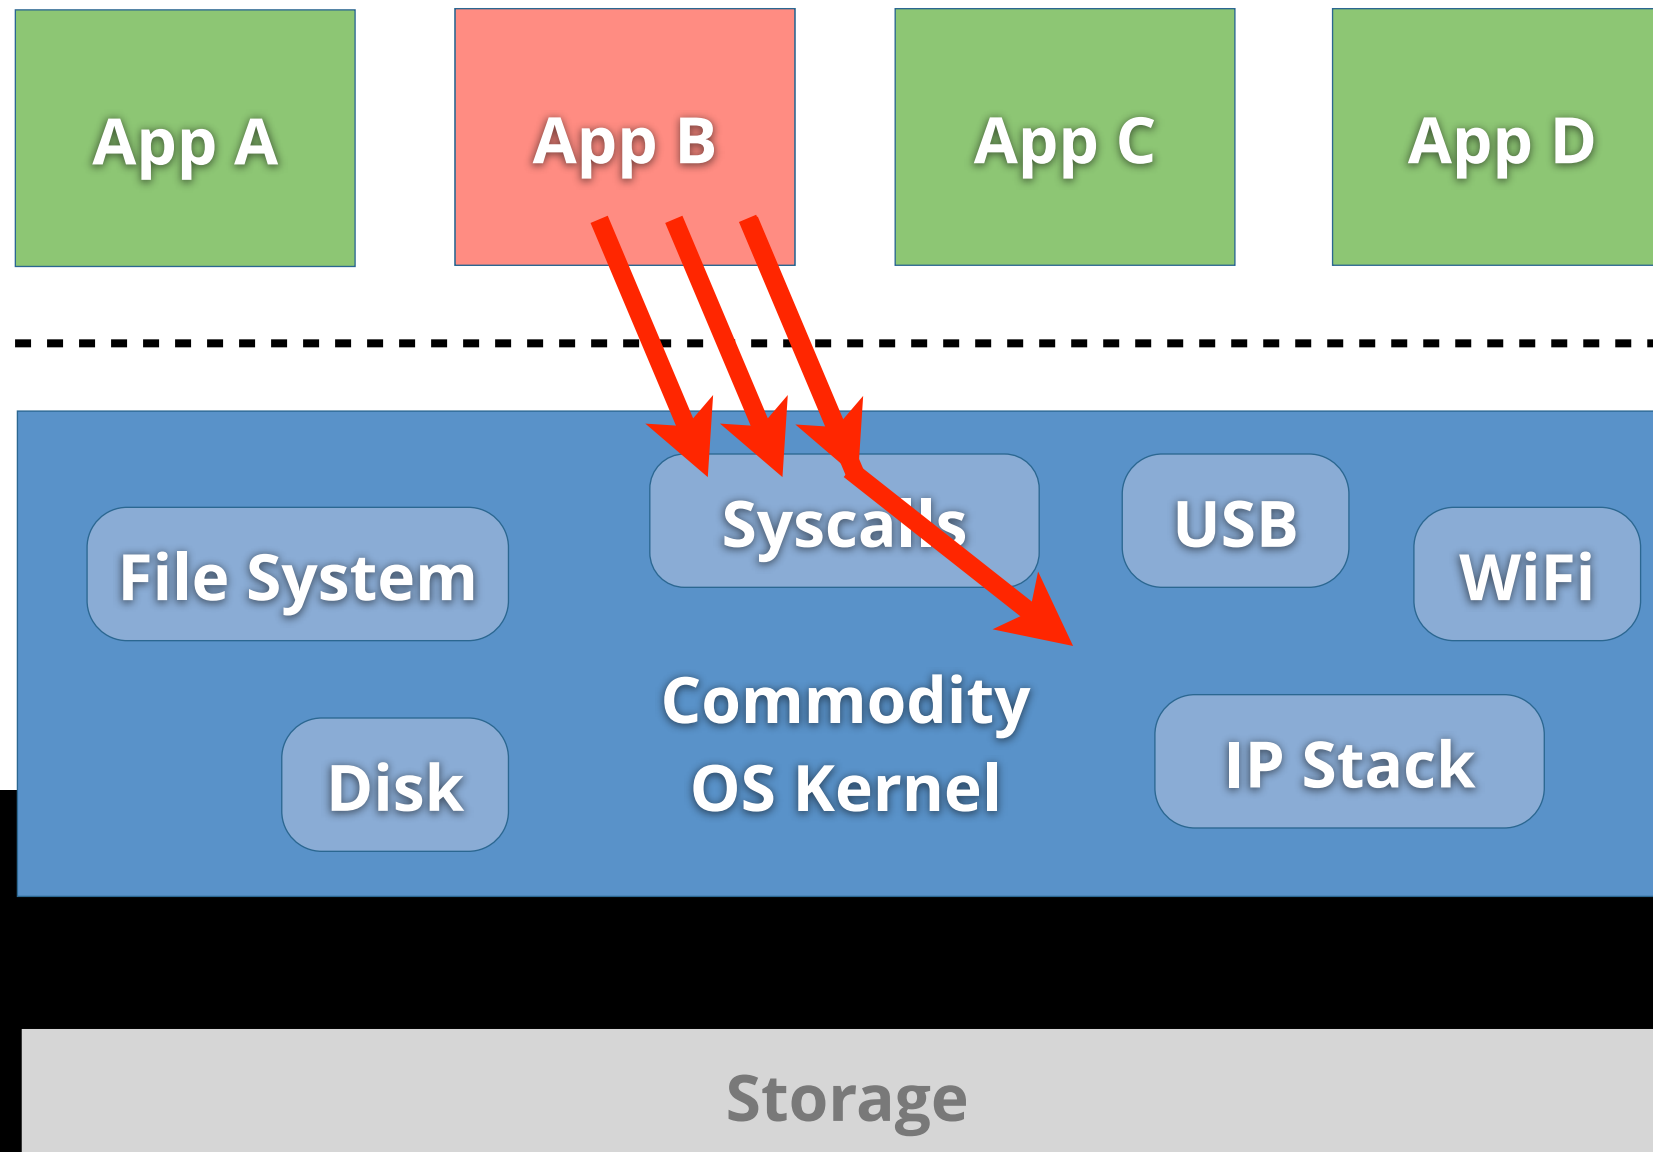
# **SECURITY ARCHITECTURES**

**CARSTEN WEINHOLD**


# CLASSICAL ARCHITECTURES

- Isolation in commodity OSes for PCs:
  - Based on user accounts
  - Same privileges for all apps
  - No isolation within applications
  - Permissive interfaces (e.g., ptrace to manipulate other address spaces)

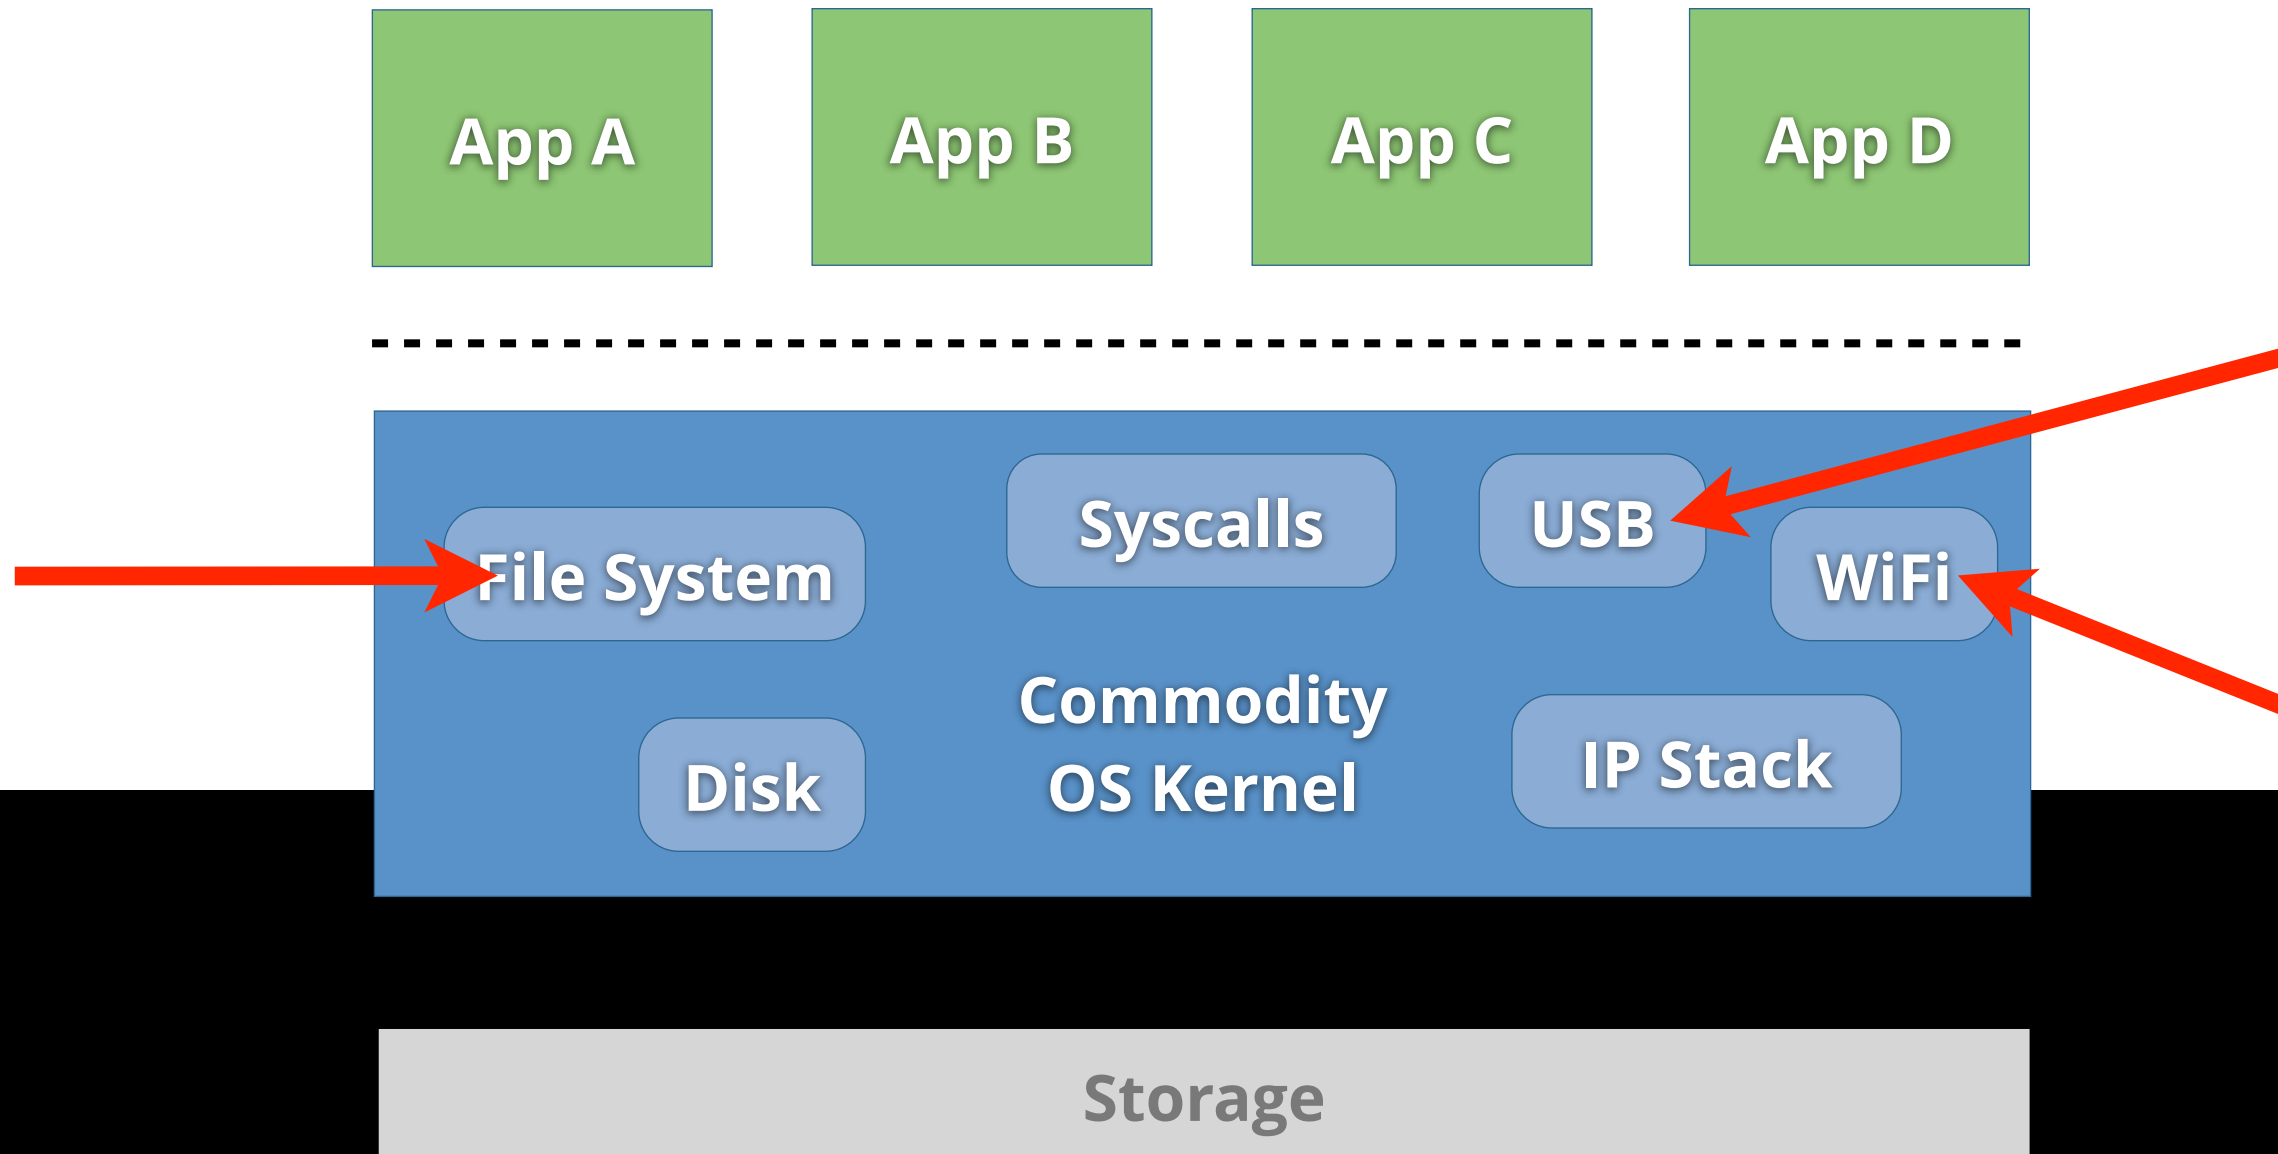
# KERNEL ATTACK VECTOR



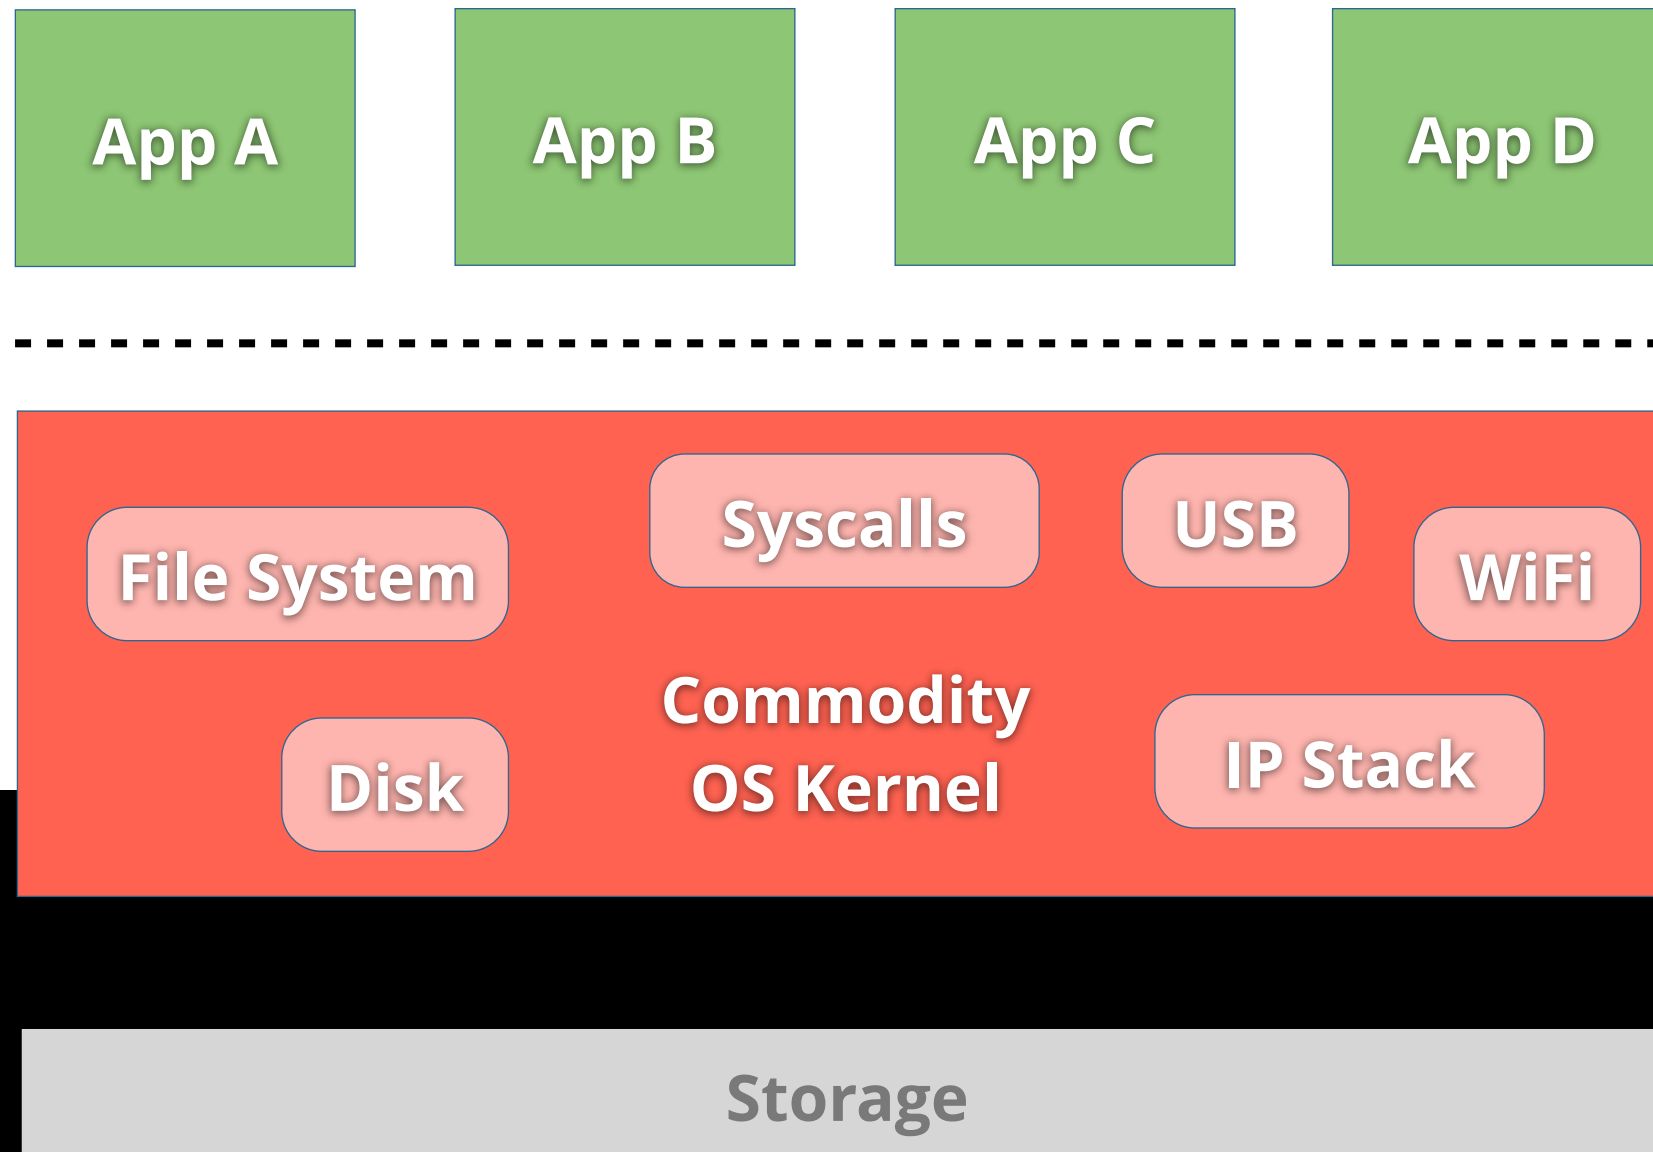
# KERNEL ATTACK VECTOR



# KERNEL ATTACK VECTOR




# KERNEL ATTACK VECTOR




- Isolation in commodity OSes for PCs:
  - Based on user accounts
  - Same privileges for all apps
  - No isolation within applications
  - Permissive interfaces (e.g., ptrace to manipulate other address spaces)



- Separate computers
- Applications and data physically isolated
- Effective, but ...
  - Higher costs



- Multiple VMs, OSes
- Isolation enforced by virtualization layer
- Saves space, energy, maintenance effort
- But still ...



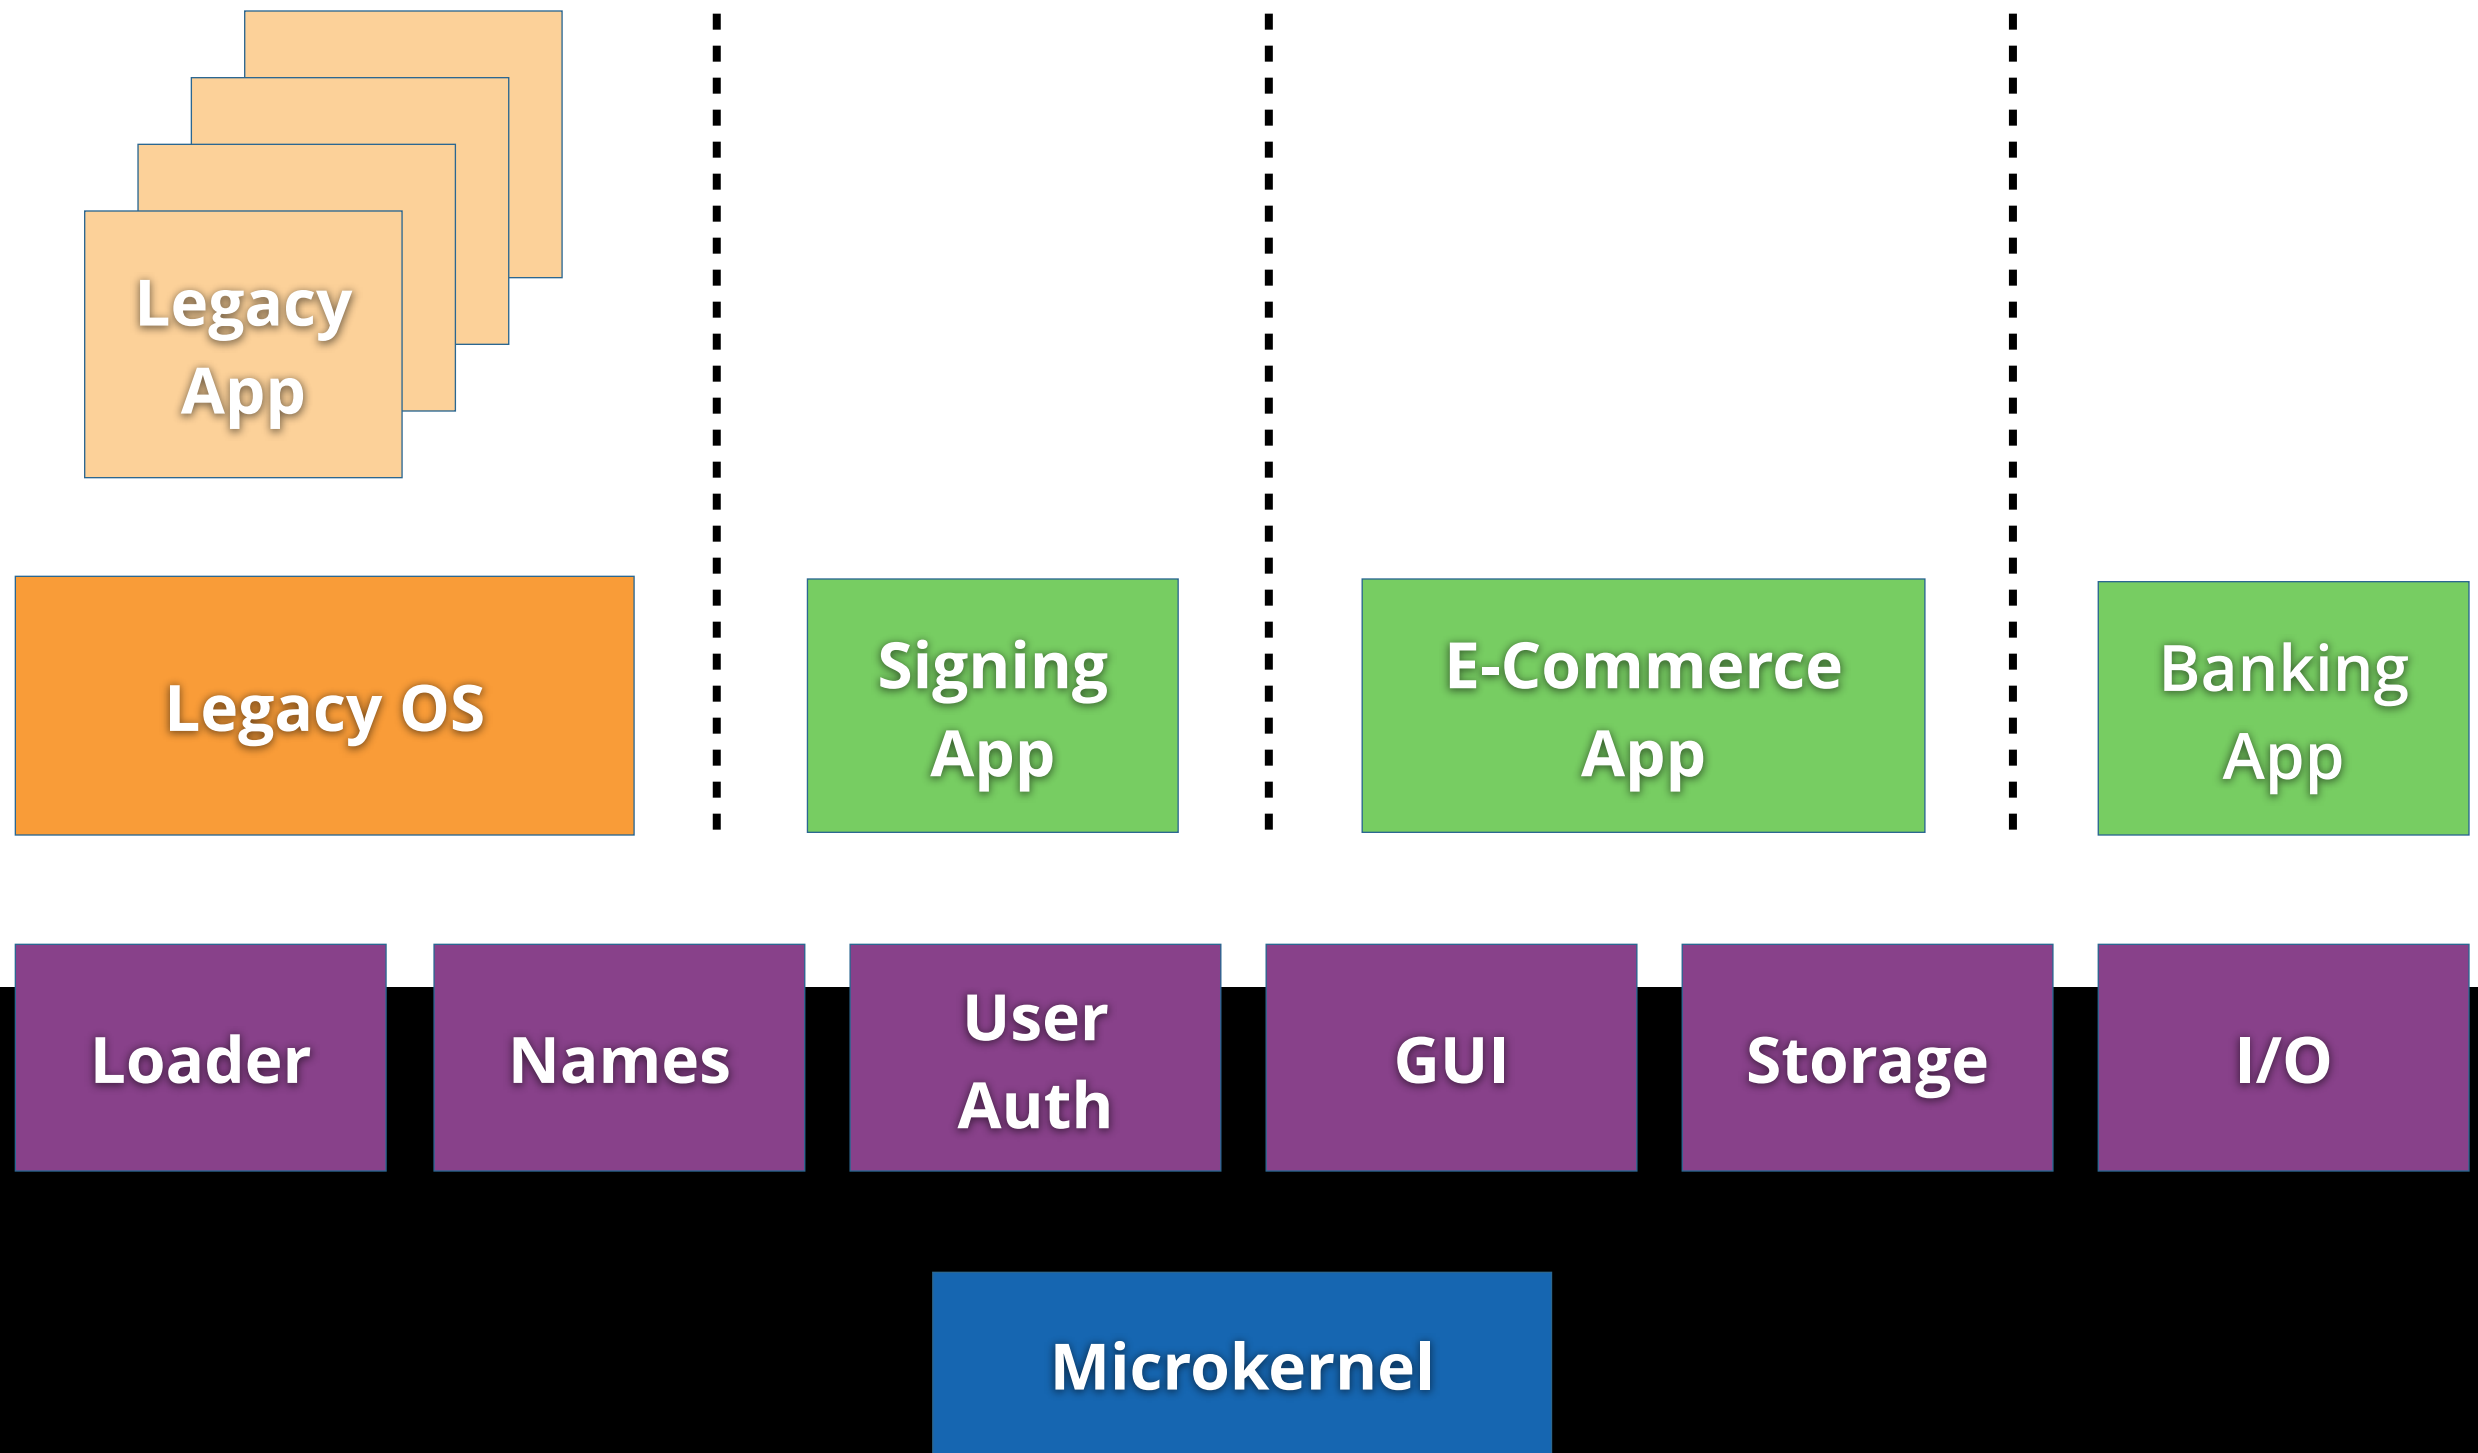
- Huge code bases remain
- Applications still the same
- Many targets to attack:
  - Applications, libraries, commodity OSes
  - Virus scanner, firewall, ...

# SECURITY ARCHITECTURES

- Protect the user's data
- Secure applications that process data
- Acknowledge different kinds of trust, e.g.:
  - Application **A** trusted to handle its own data, but not the files of application **B**

- To improve security: Reduce size of TCB  
= smaller attack surface
- First (incomplete) idea:
  - Remove huge legacy OS from TCB
  - Port application to microkernel-based multi-server OS

# NIZZA ARCHITECTURE



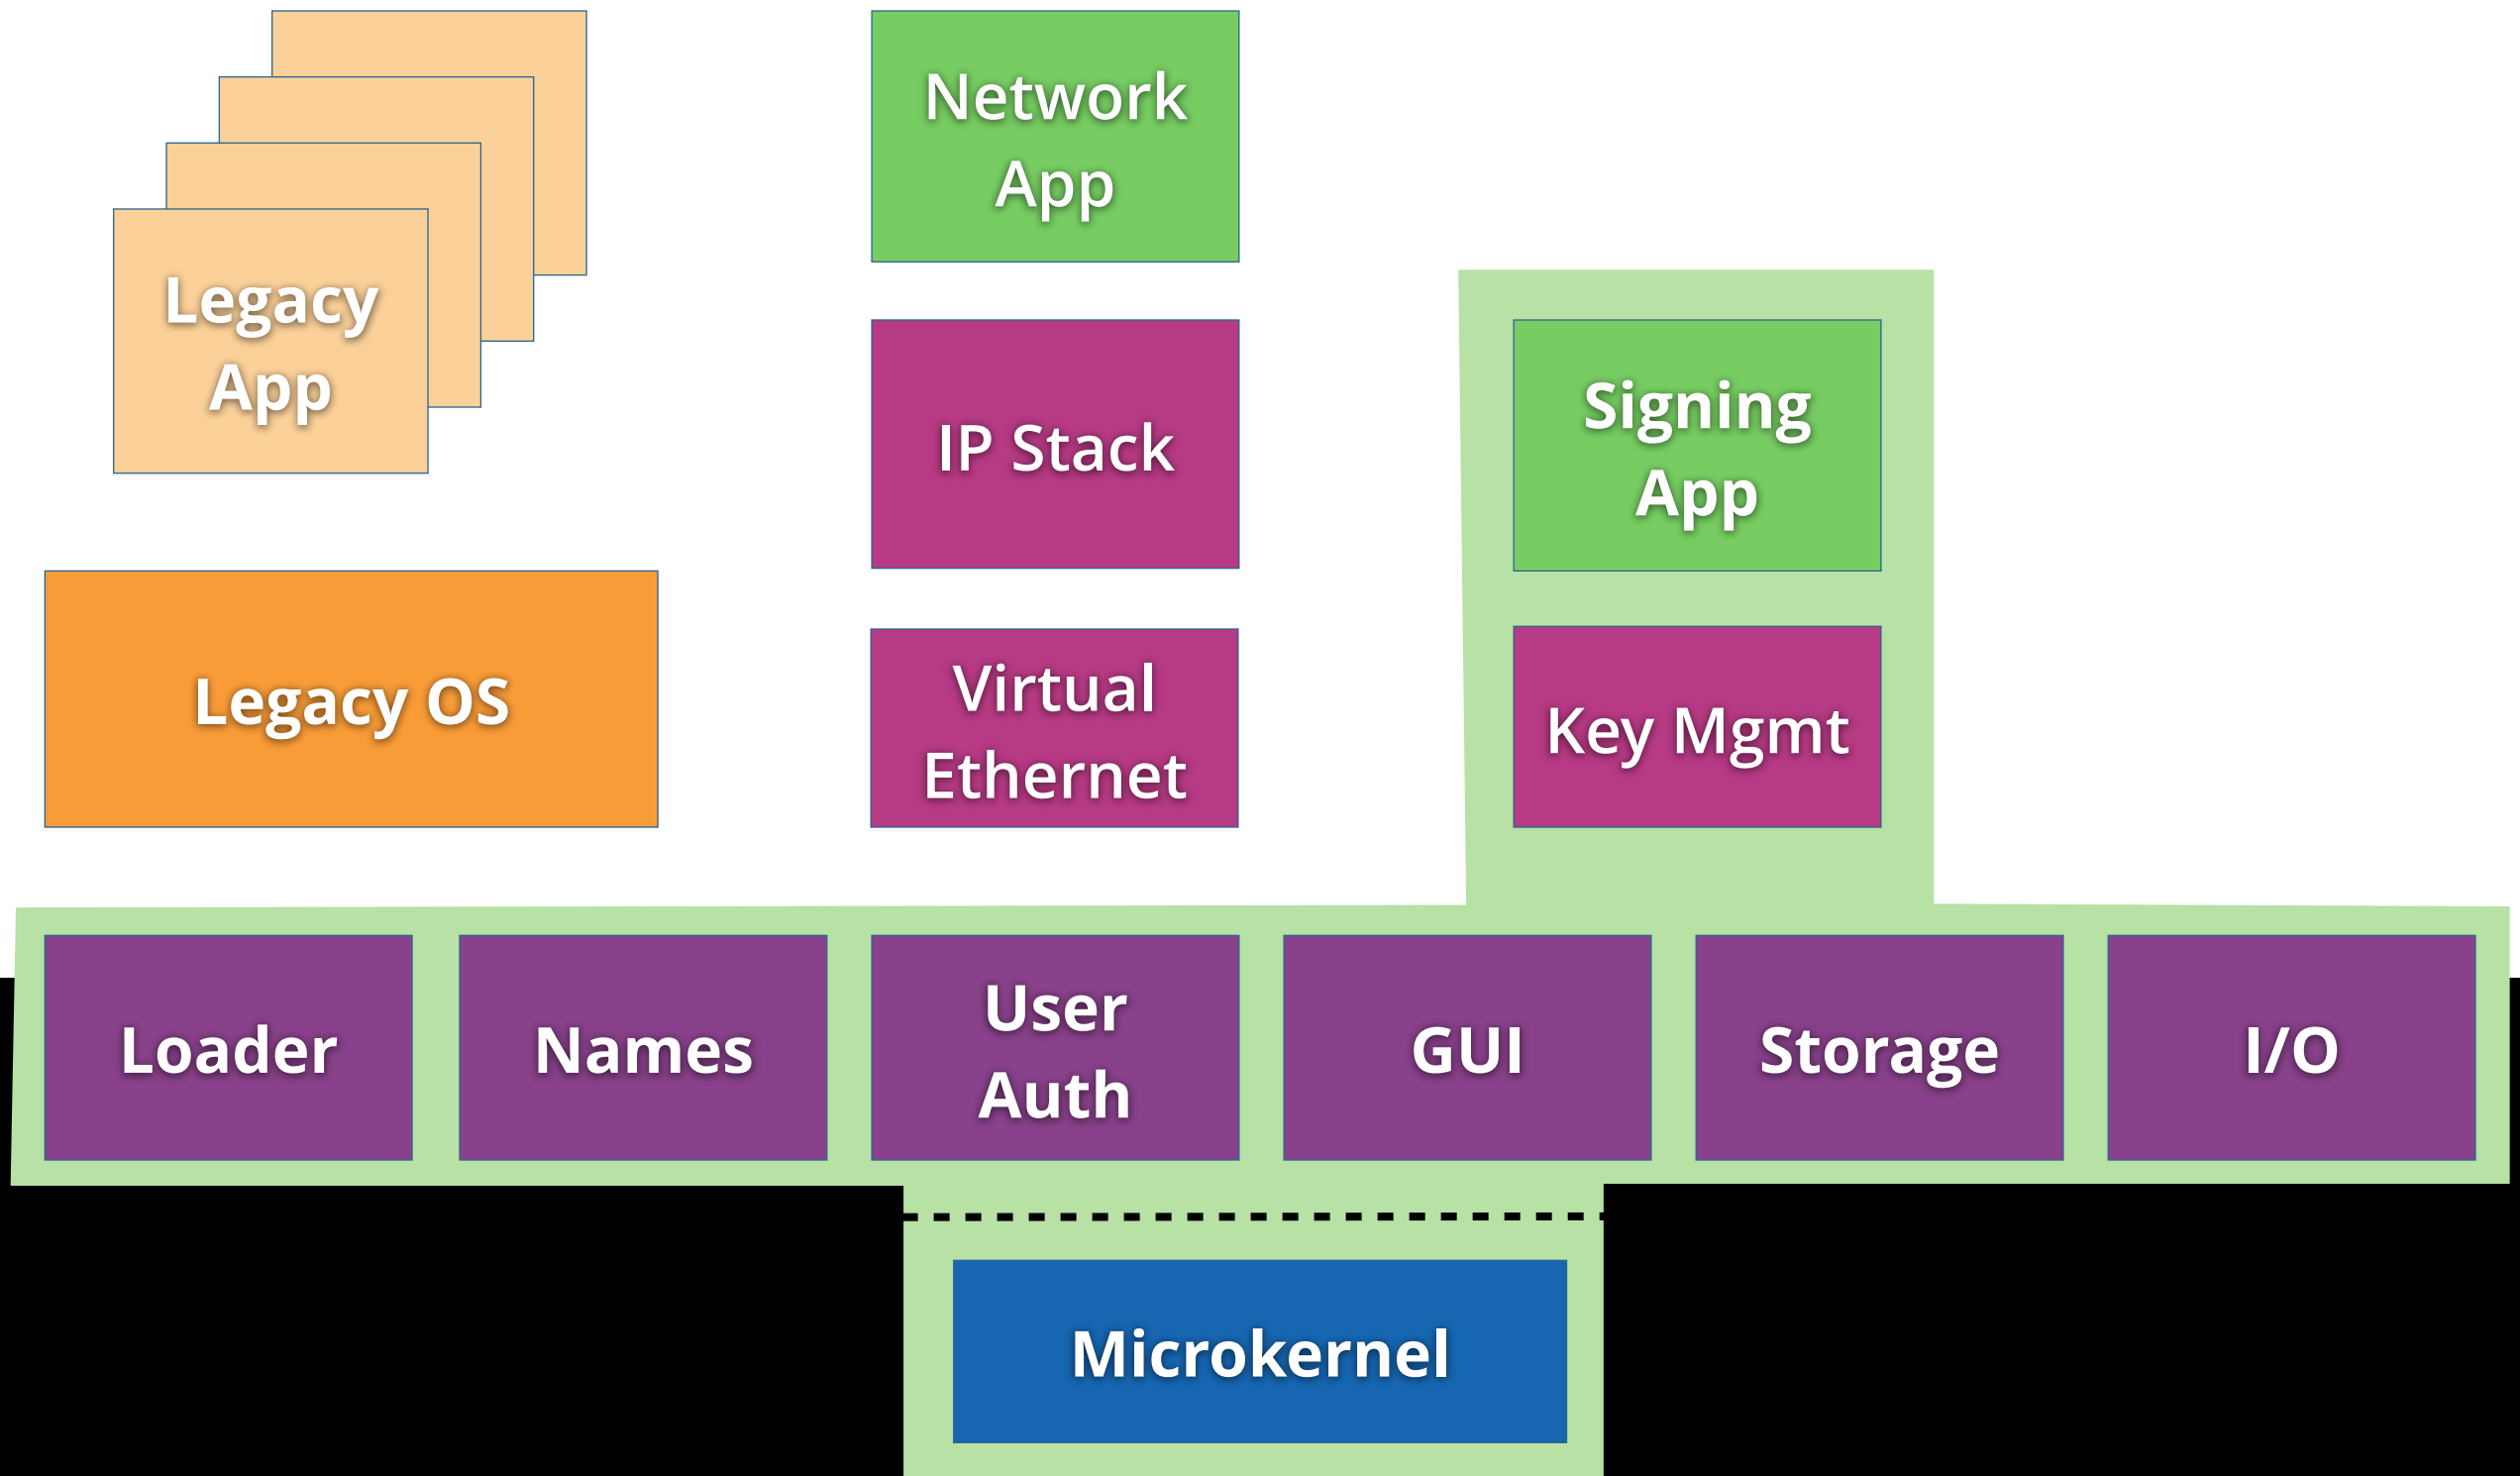
## **Nizza architecture: fundamental concepts:**

- Strong isolation
  - Application-specific TCBs
  - Legacy reuse
- 
- Trusted computing



- Reflects **Principle of Least Privilege**
- TCB of an application includes only components its security relies upon
- TCB does not include unrelated applications, services, libraries

# APP-SPECIFIC TCB



- Reflects **Principle of Least Privilege**
- TCB of an application includes only components its security relies upon
- TCB does not include unrelated applications, services, libraries
- **Mechanisms:**



# SPLITTING COMPONENTS

- Problems with porting applications:
  - Dependencies need to be satisfied
  - Can be complex, require lots of code
  - Stripped down applications may lack functionality / usability
- Better idea: split application

## Digitally signed e-mails, what's critical?

- Handling of signature keys
- Requesting passphrase to unlock signature key
- Presenting e-mail message:
  - Before sending: **What You See Is What**

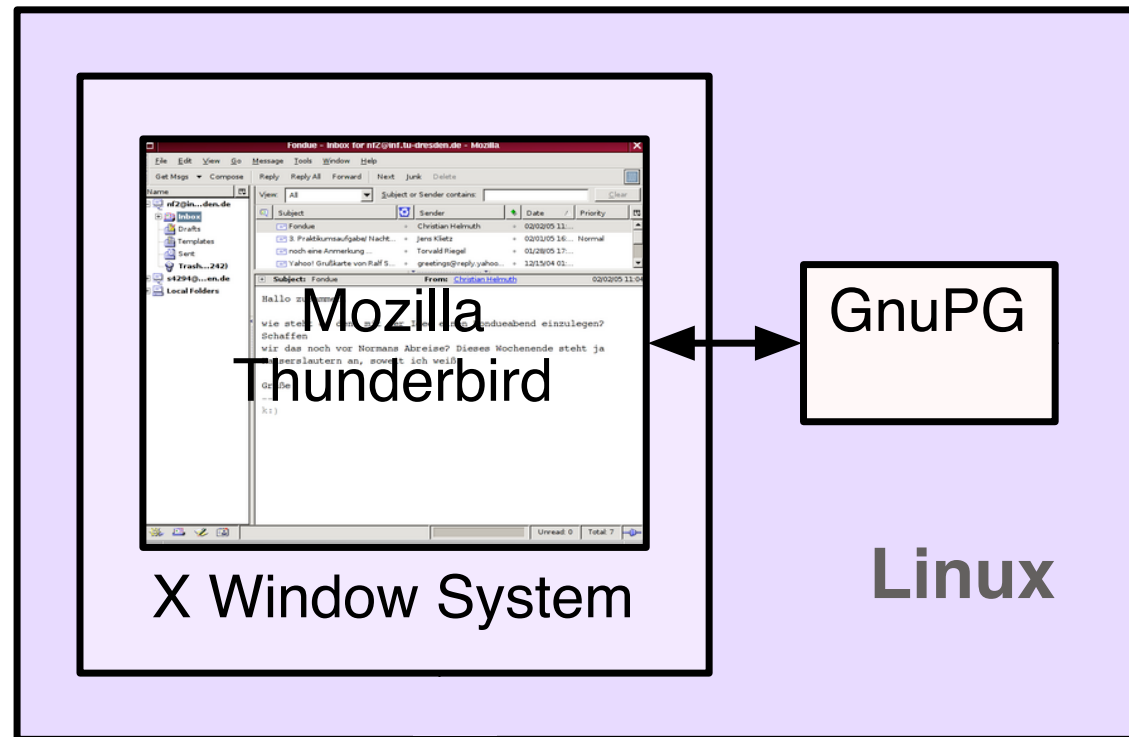


Image source: [5]

# SPLIT EMAIL APP

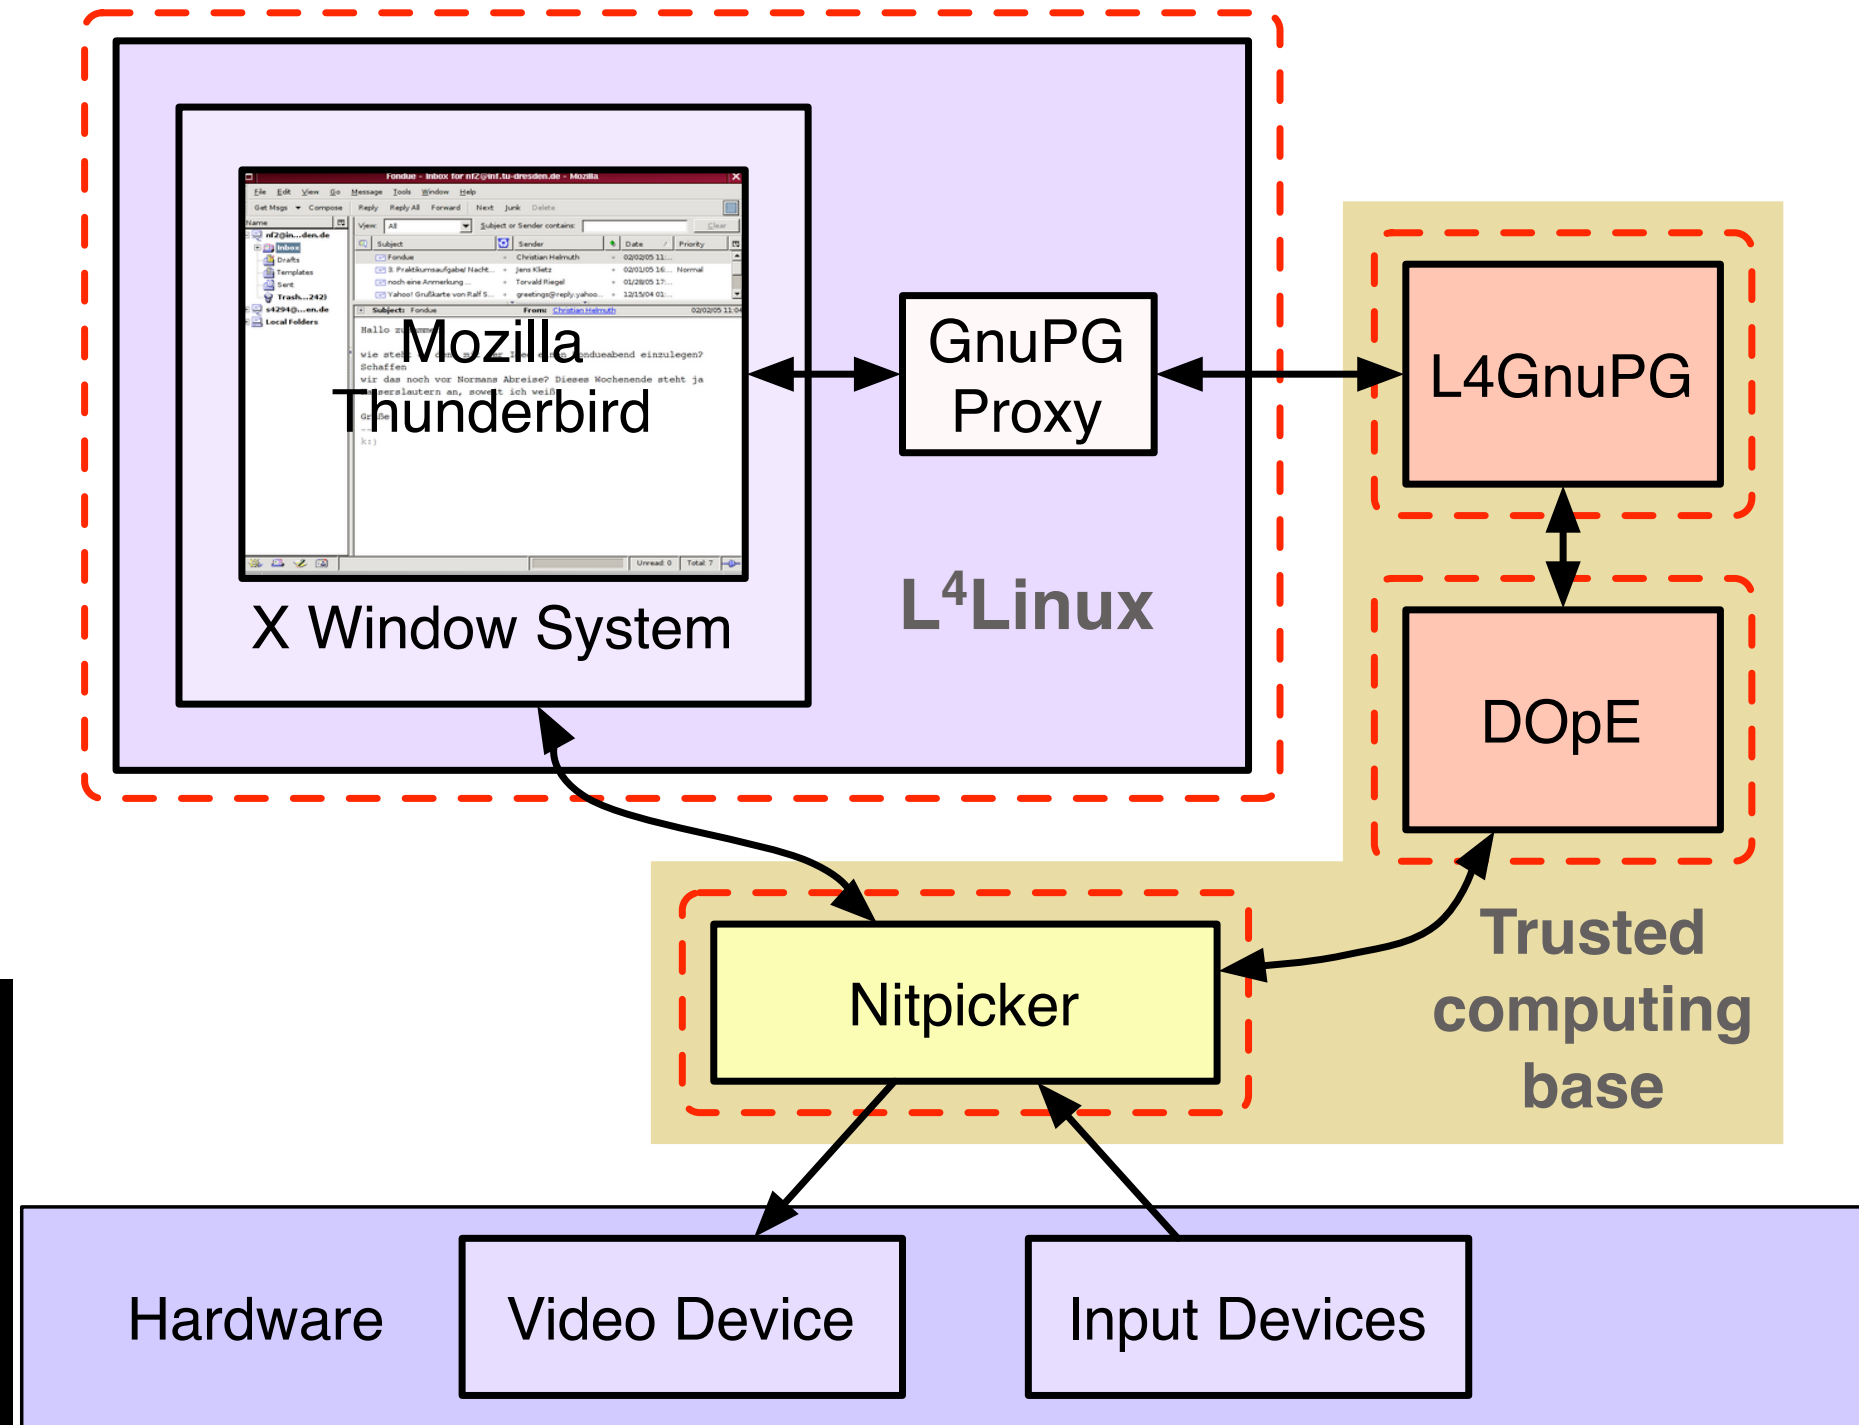


Image source: [5]



- *1,500,000+* SLOC no longer in TCB:
  - Linux kernel, drivers, X-Server
  - C and GUI libraries, Thunderbird, ...
- TCB size reduced to *~150,000* SLOC:
  - GNU Privacy Guard, e-mail viewer

- Splitting works for applications
- What about the complex and useful infrastructure of commodity OSes?
  - Drivers (see previous lectures)
  - Protocol stacks (e.g., TCP/IP)


- Run legacy OS in VM
- Reuse service: net, files, ...
- Legacy infrastructure isolated from applications
- But:
  - Applications still depend on



- Network and file system stacks are virtually essential subsystems
- Generally well tested
- Ready for production use
- ... but not bug free [1,2]:
  - Linux file systems (LFS, ISO 9660, Ext3)

- Complex protocol stacks should not be part of TCB (for confidentiality + integrity)
- Reuse untrusted infrastructure through **Trusted Wrapper:**
  - Add security around existing APIs
    - Cryptography

**VPN:** Confidentiality, Integrity, ~~Availability~~




- SINA box used by German „BSI“:
- VPN gateway
- Implements IPSec & PKI
- Intrusion detection & response



Image source:

<http://www.secunet.com/de/das-unternehmen/presse/bilddatenbank/>

- Differently trusted network interfaces:
  - **Red:** plaintext, no protection
  - **Black:** encryption + authentication codes





- Linux is complex!
  - SLOC for Linux 2.6.18:
    - Architecture specific: 817,880
    - x86 specific: 55,463
    - Drivers: 2,365,256
    - Core: 4,000,507
- Released date:  
20 Sep 2006**

- Linux is even more complex in 2024!
- SLOC for Linux 6.7.1:
  - Architecture specific: 1,729,519
  - x86 specific: 316,544
  - Drivers: 17,771,667
  - Core utilities: 4,062,225

- Research project „Mikro-SINA“
- Goals:
  - Reduce TCB of VPN gateway software
  - Enable high-level evaluation for high assurance scenarios

- Protocol suite for securing IP-based communication
- Authentication header (**AH**)
  - Integrity
  - Authentication
- Encapsulating Security Payload

Application

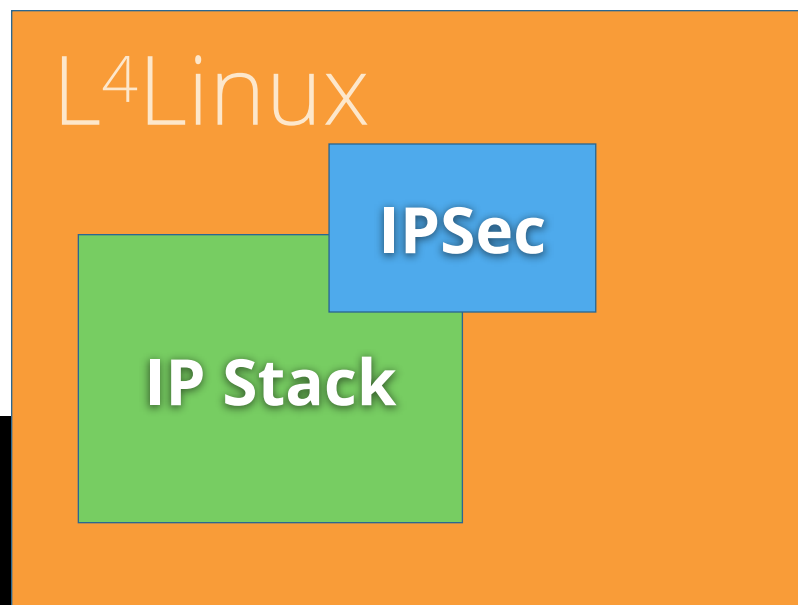
TCP / UDP

IP


IPSec

Link Layer

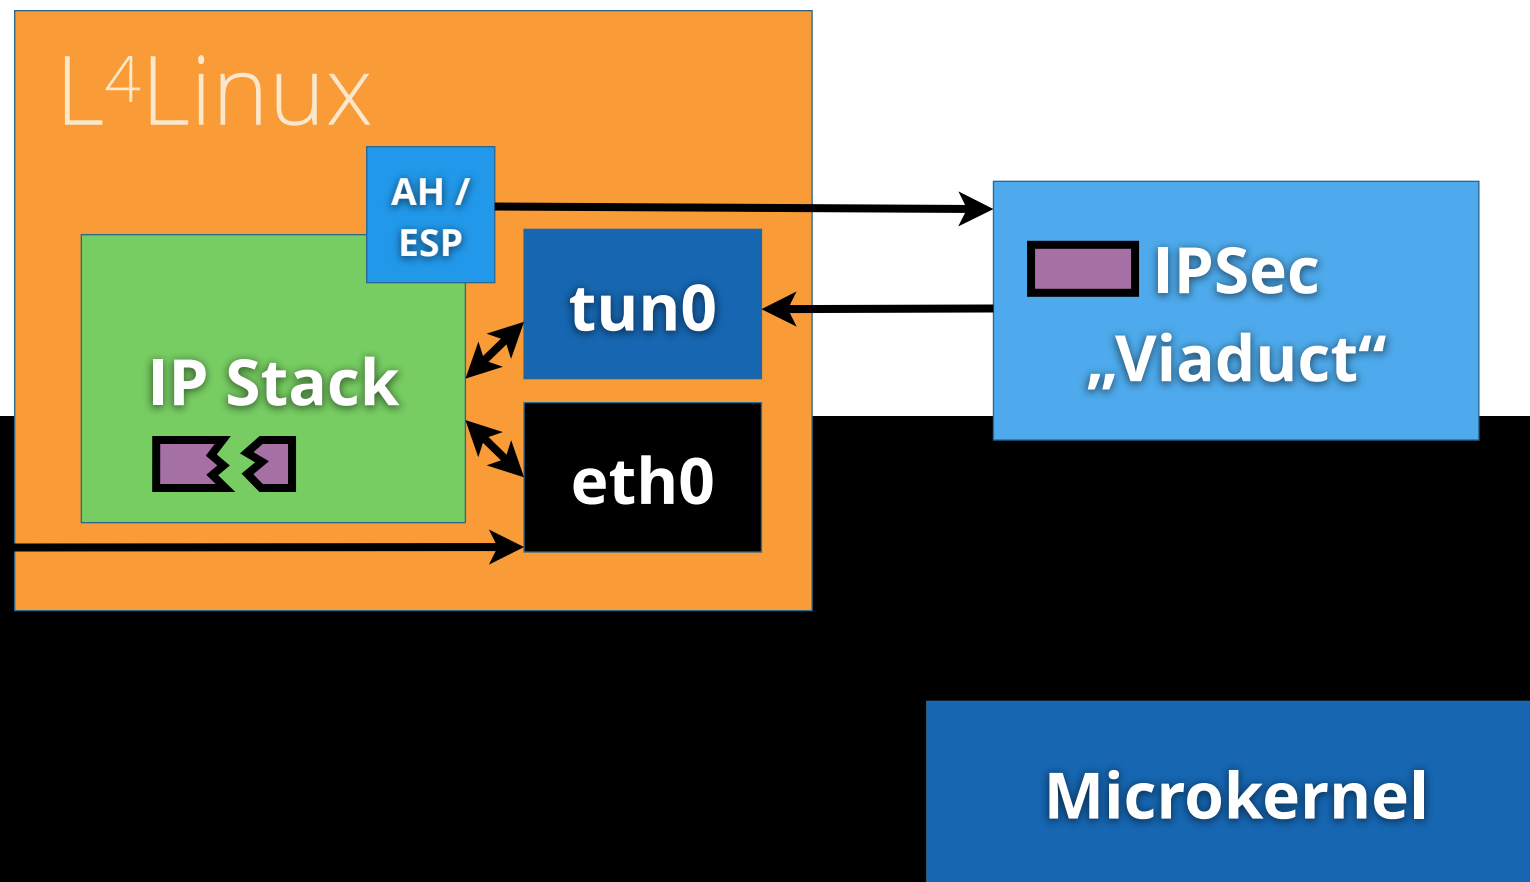
- IPsec is security critical component
- ... but is integrated into Linux kernel



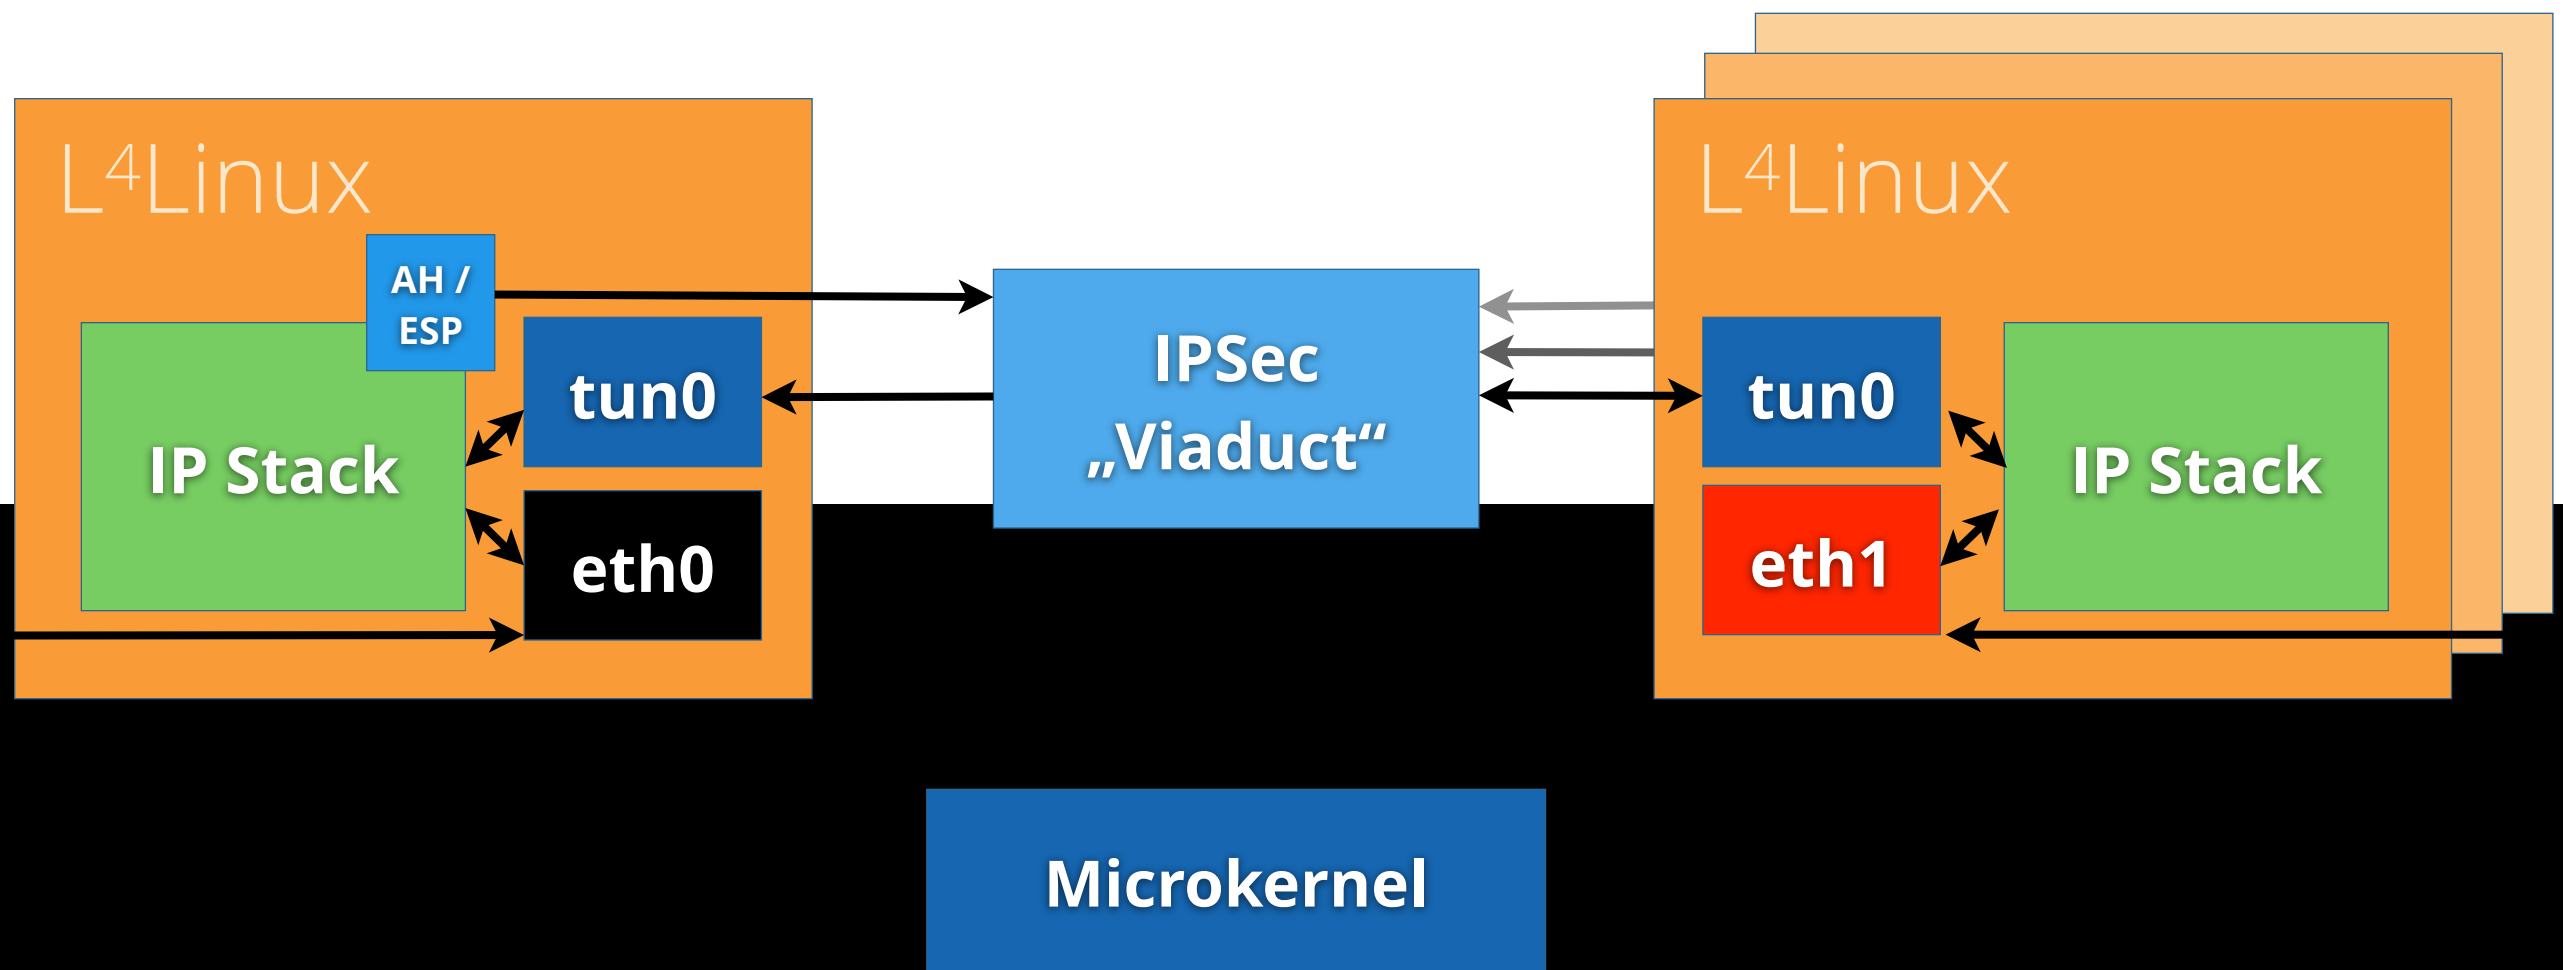
- **Idea:** Isolate IPsec in „Viaduct“
- IPsec packets sent/received through TUN/TAP device



- Problem: Routers can fragment IPSec packets on the way
- Let L<sup>4</sup>Linux reassemble them



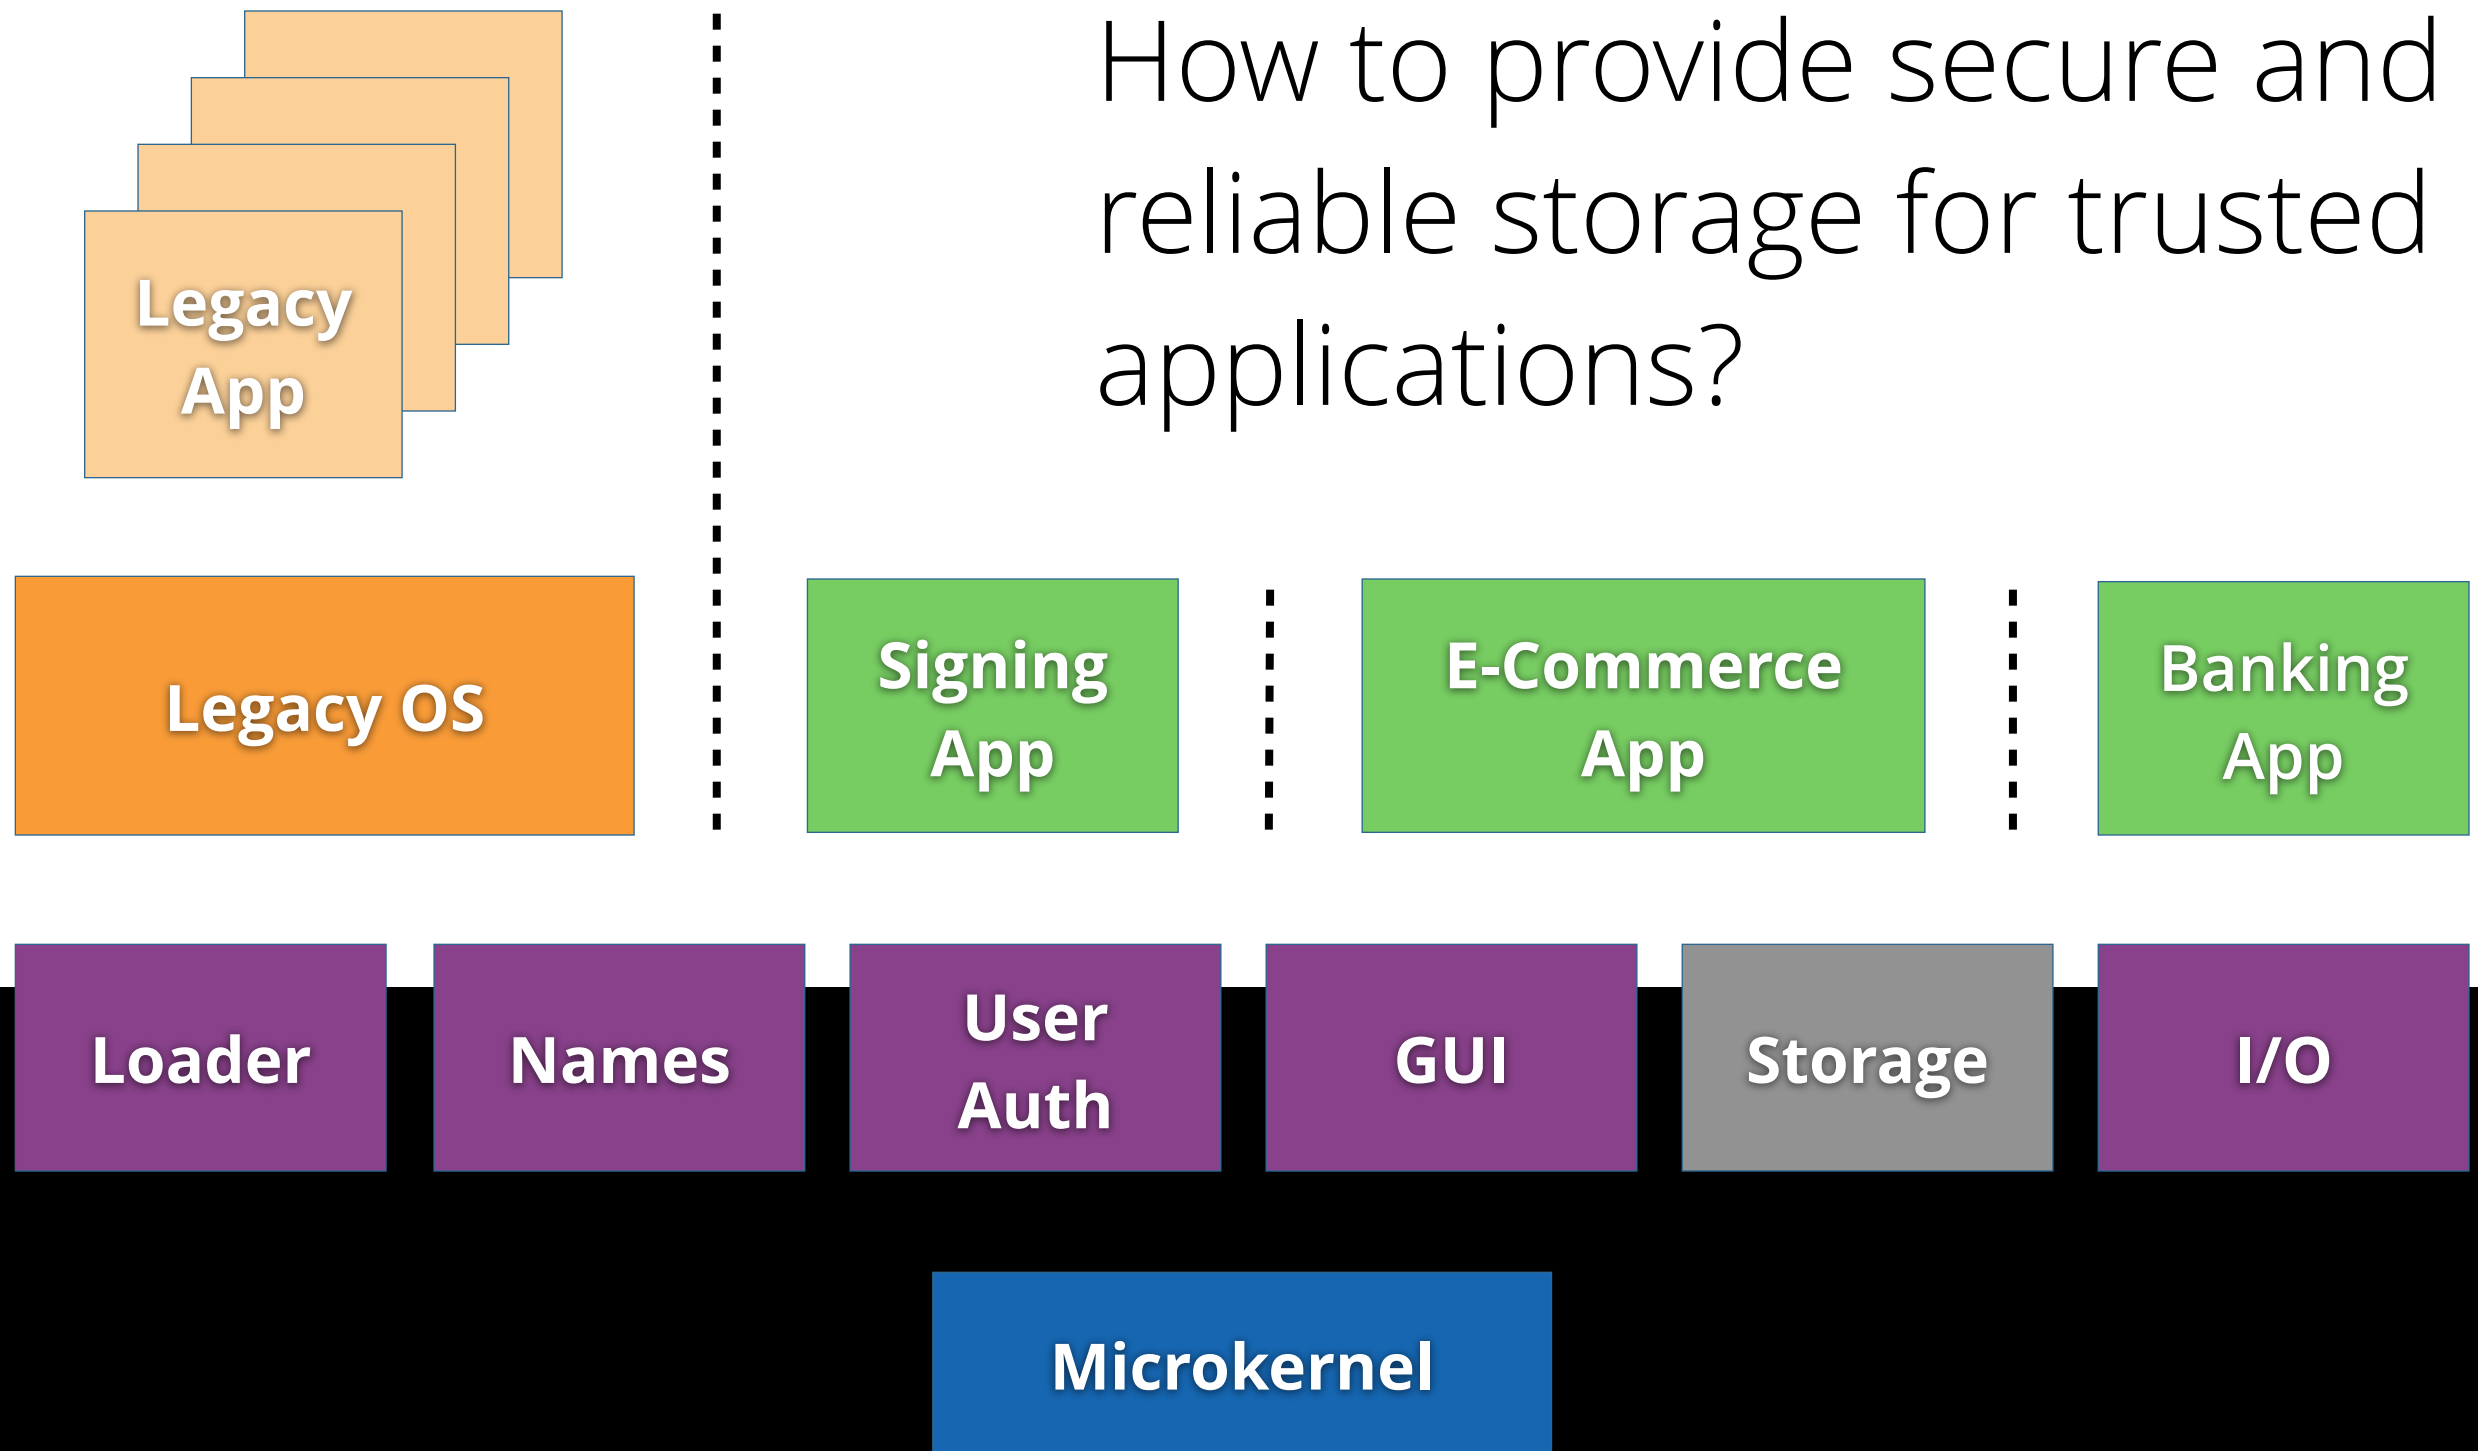
- Untrusted L<sup>4</sup>Linux instances must not see both plaintext and encrypted data
- Dedicated L<sup>4</sup>Linux for black/red networks



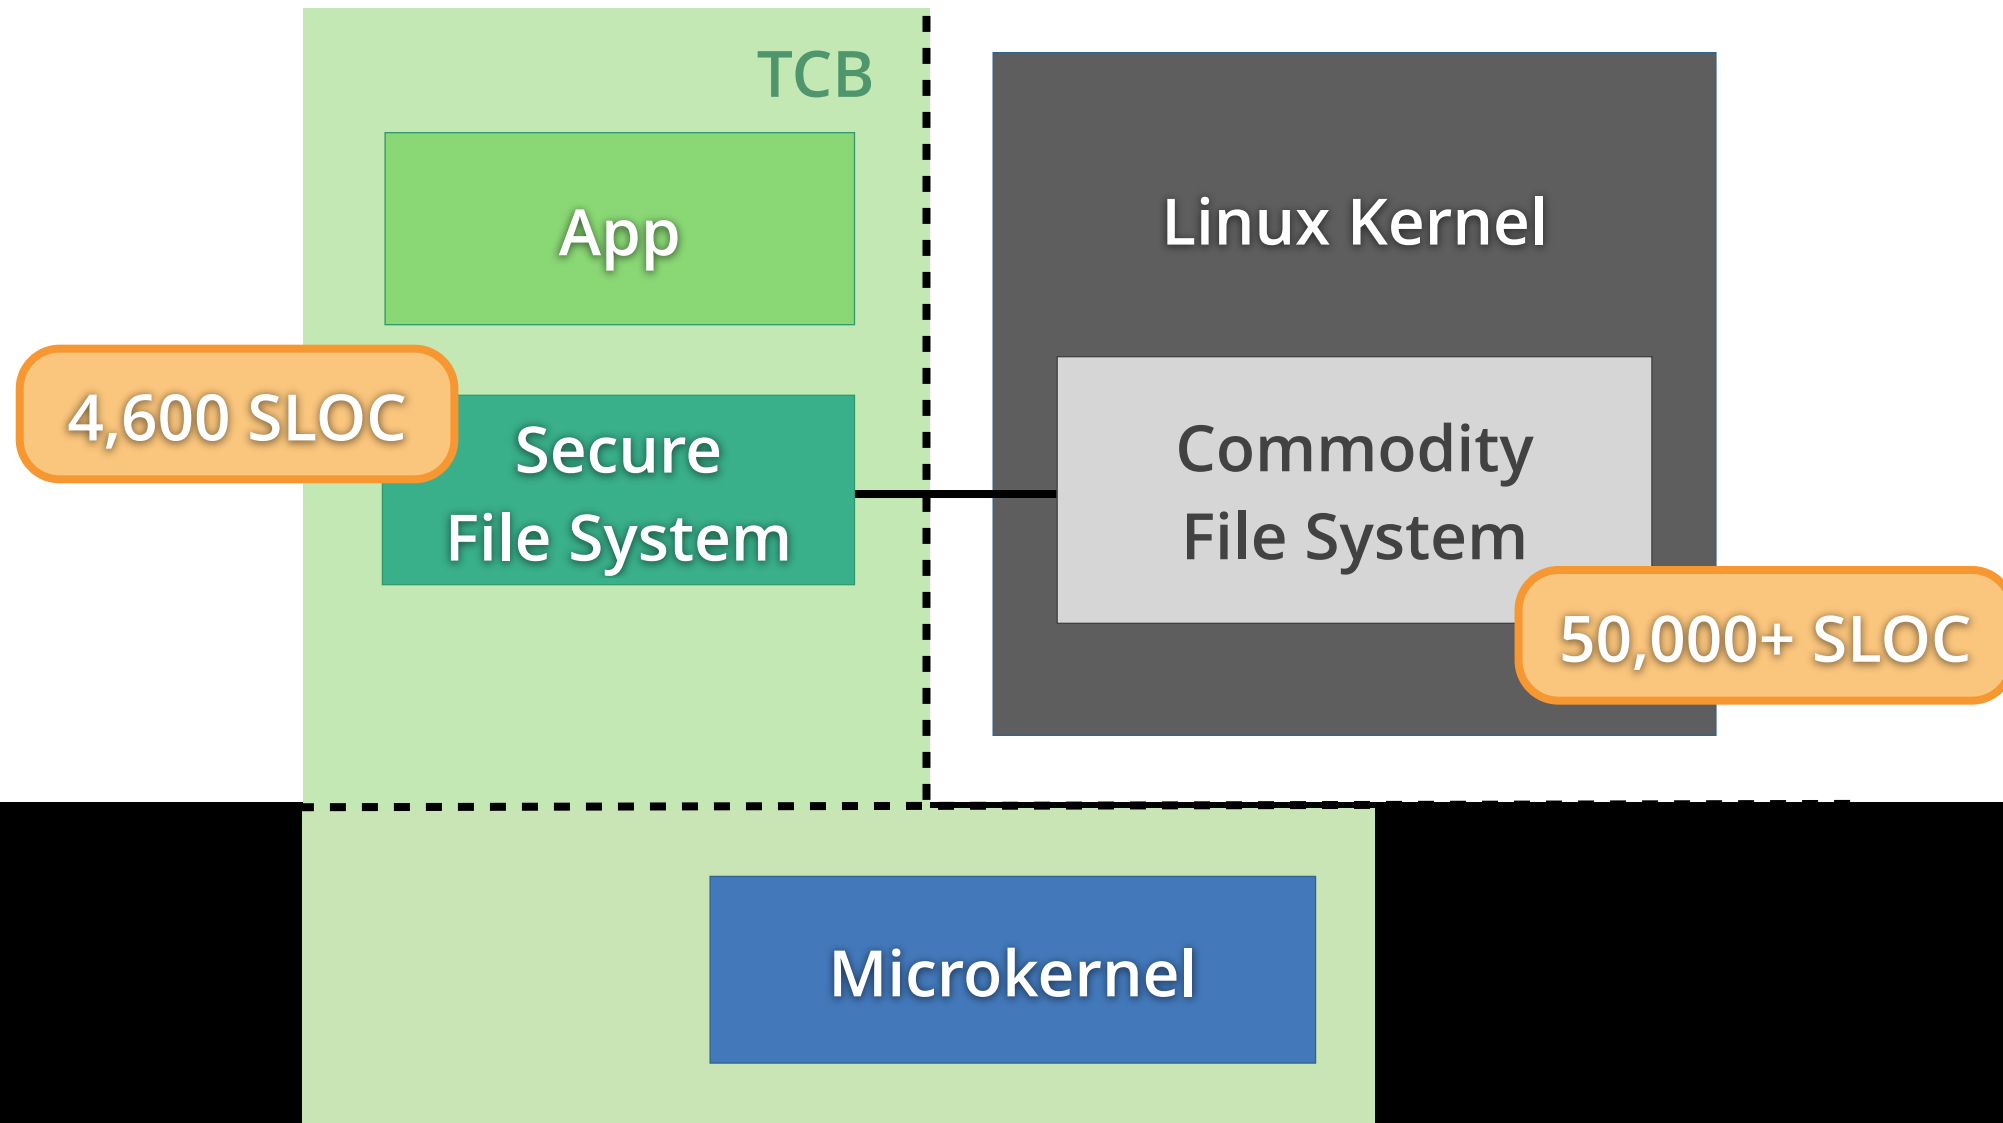


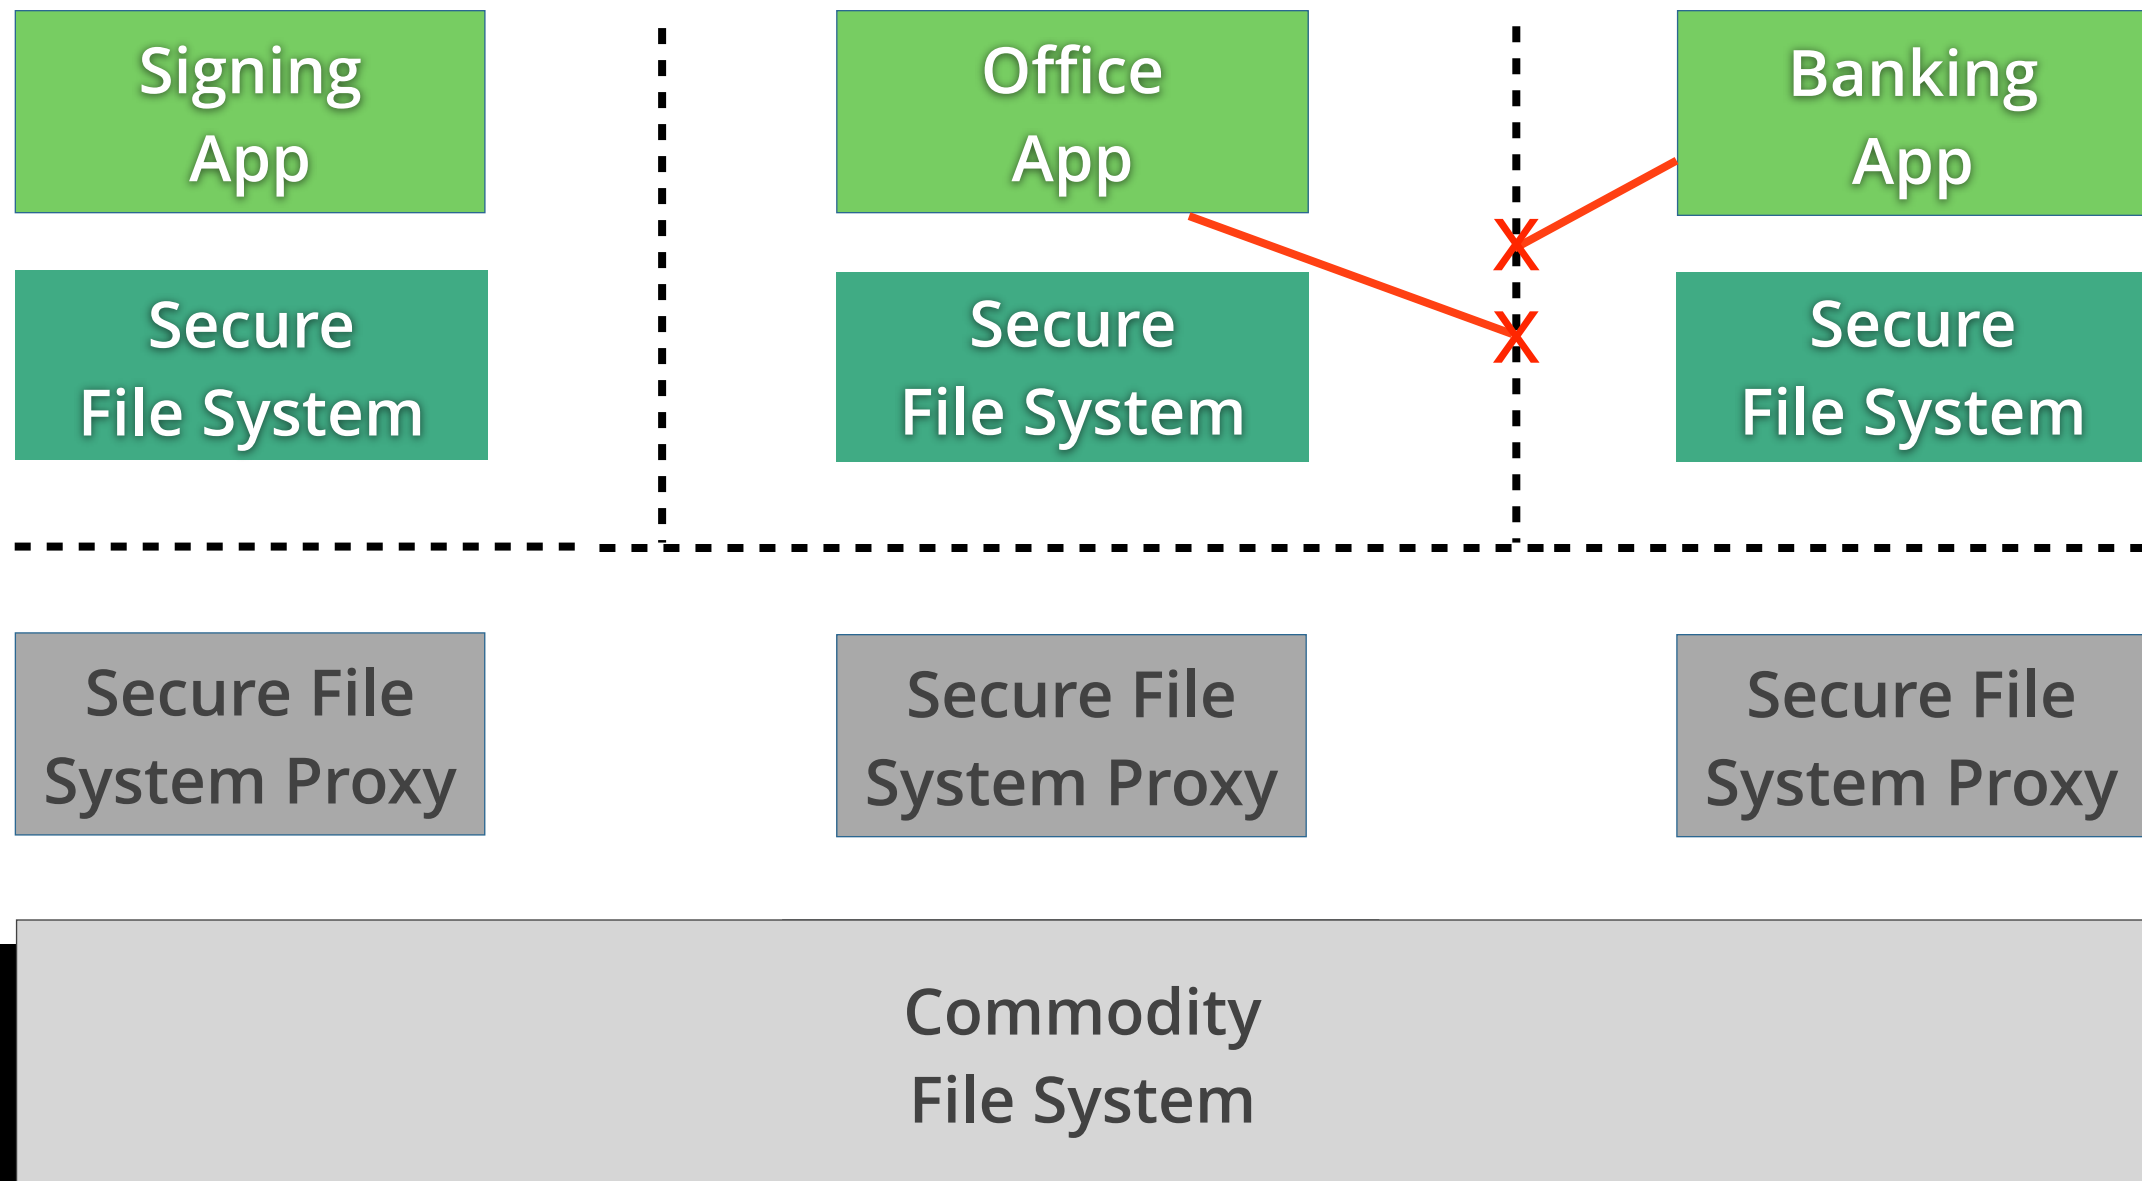
- Result: Trusted Wrapper for VPN
- Small TCB (see [6] for details):
  - 5,000 SLOC for „Viaduct“
  - Fine grain isolation
  - Principle of least privilege

# EXAMPLE 3: STORAGE

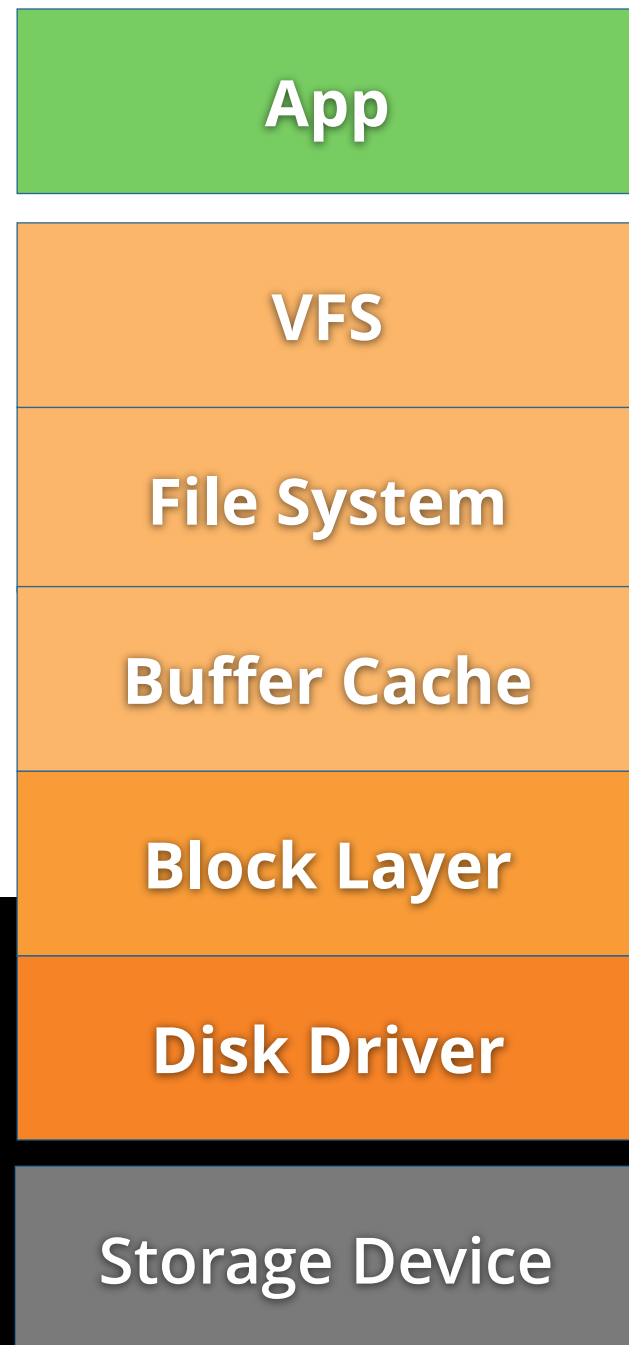


**VPFS:** Confidentiality, Integrity, ~~Availability~~





- **Confidentiality:** only authorized applications can access file system, all untrusted software cannot get any useful information
- **Integrity:** all data and meta data is correct, complete, and up to date;
- **Recoverability:** damaged data in untrusted file system can be recovered



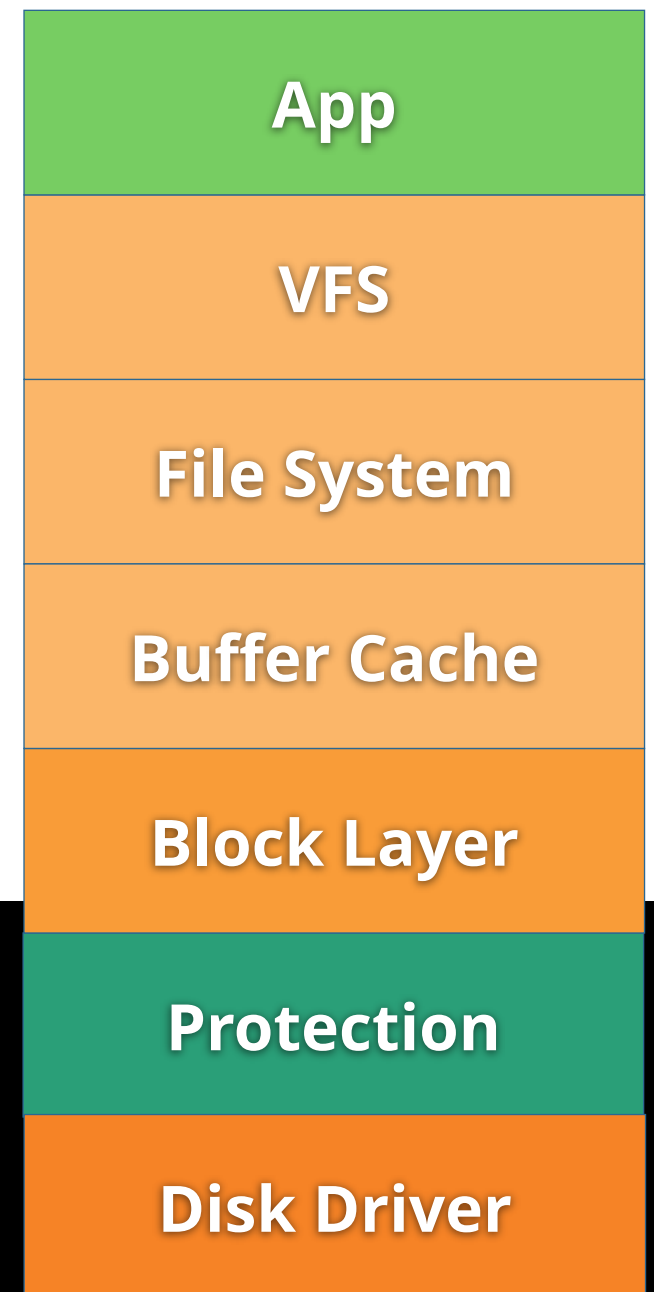
## File-level protection

- CFS** Cryptographic File System for UNIX
- EFS** Microsoft Encrypting File System
- ecryptfs** Linux kernel support + tools
- EncFS** Based on FUSE

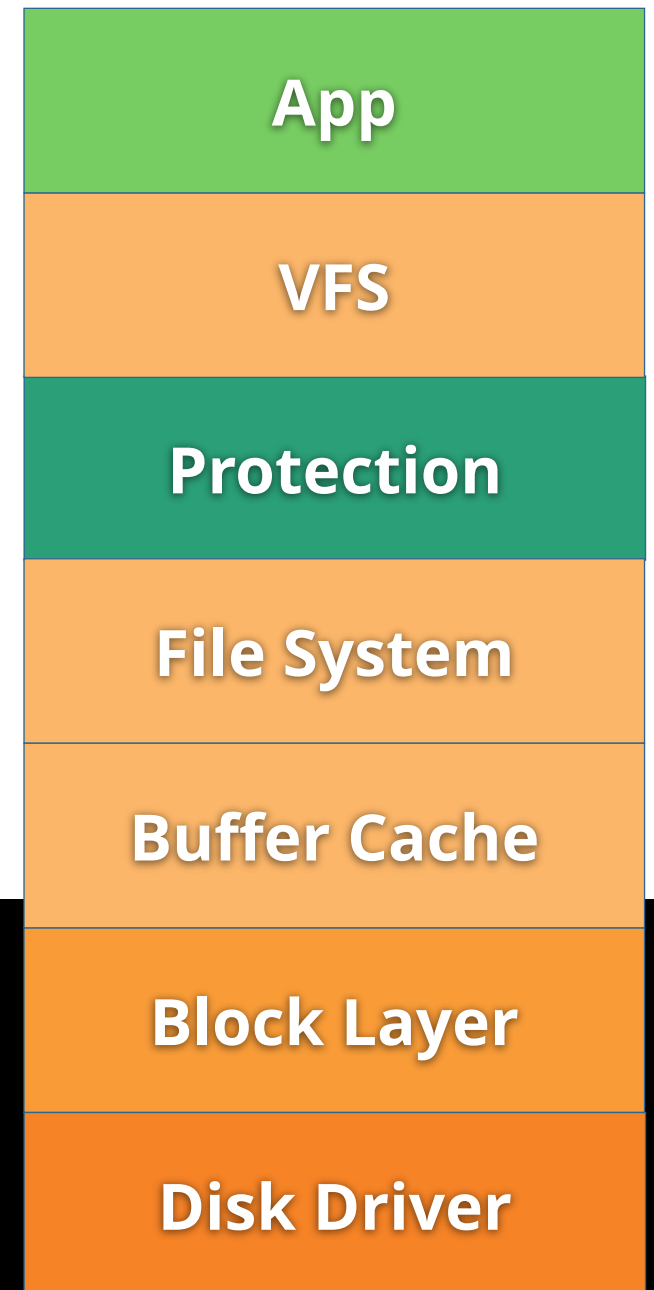
## Volume-level protection

- TrueCrypt, Filevault 2**
- dm\_crypt**
- Bitlocker**
- Encrypted volumes in smartphones, etc.

- First end of design space:  
Protect at block layer
- Transparent encryption of all  
data and metadata
- Block-level integrity ???

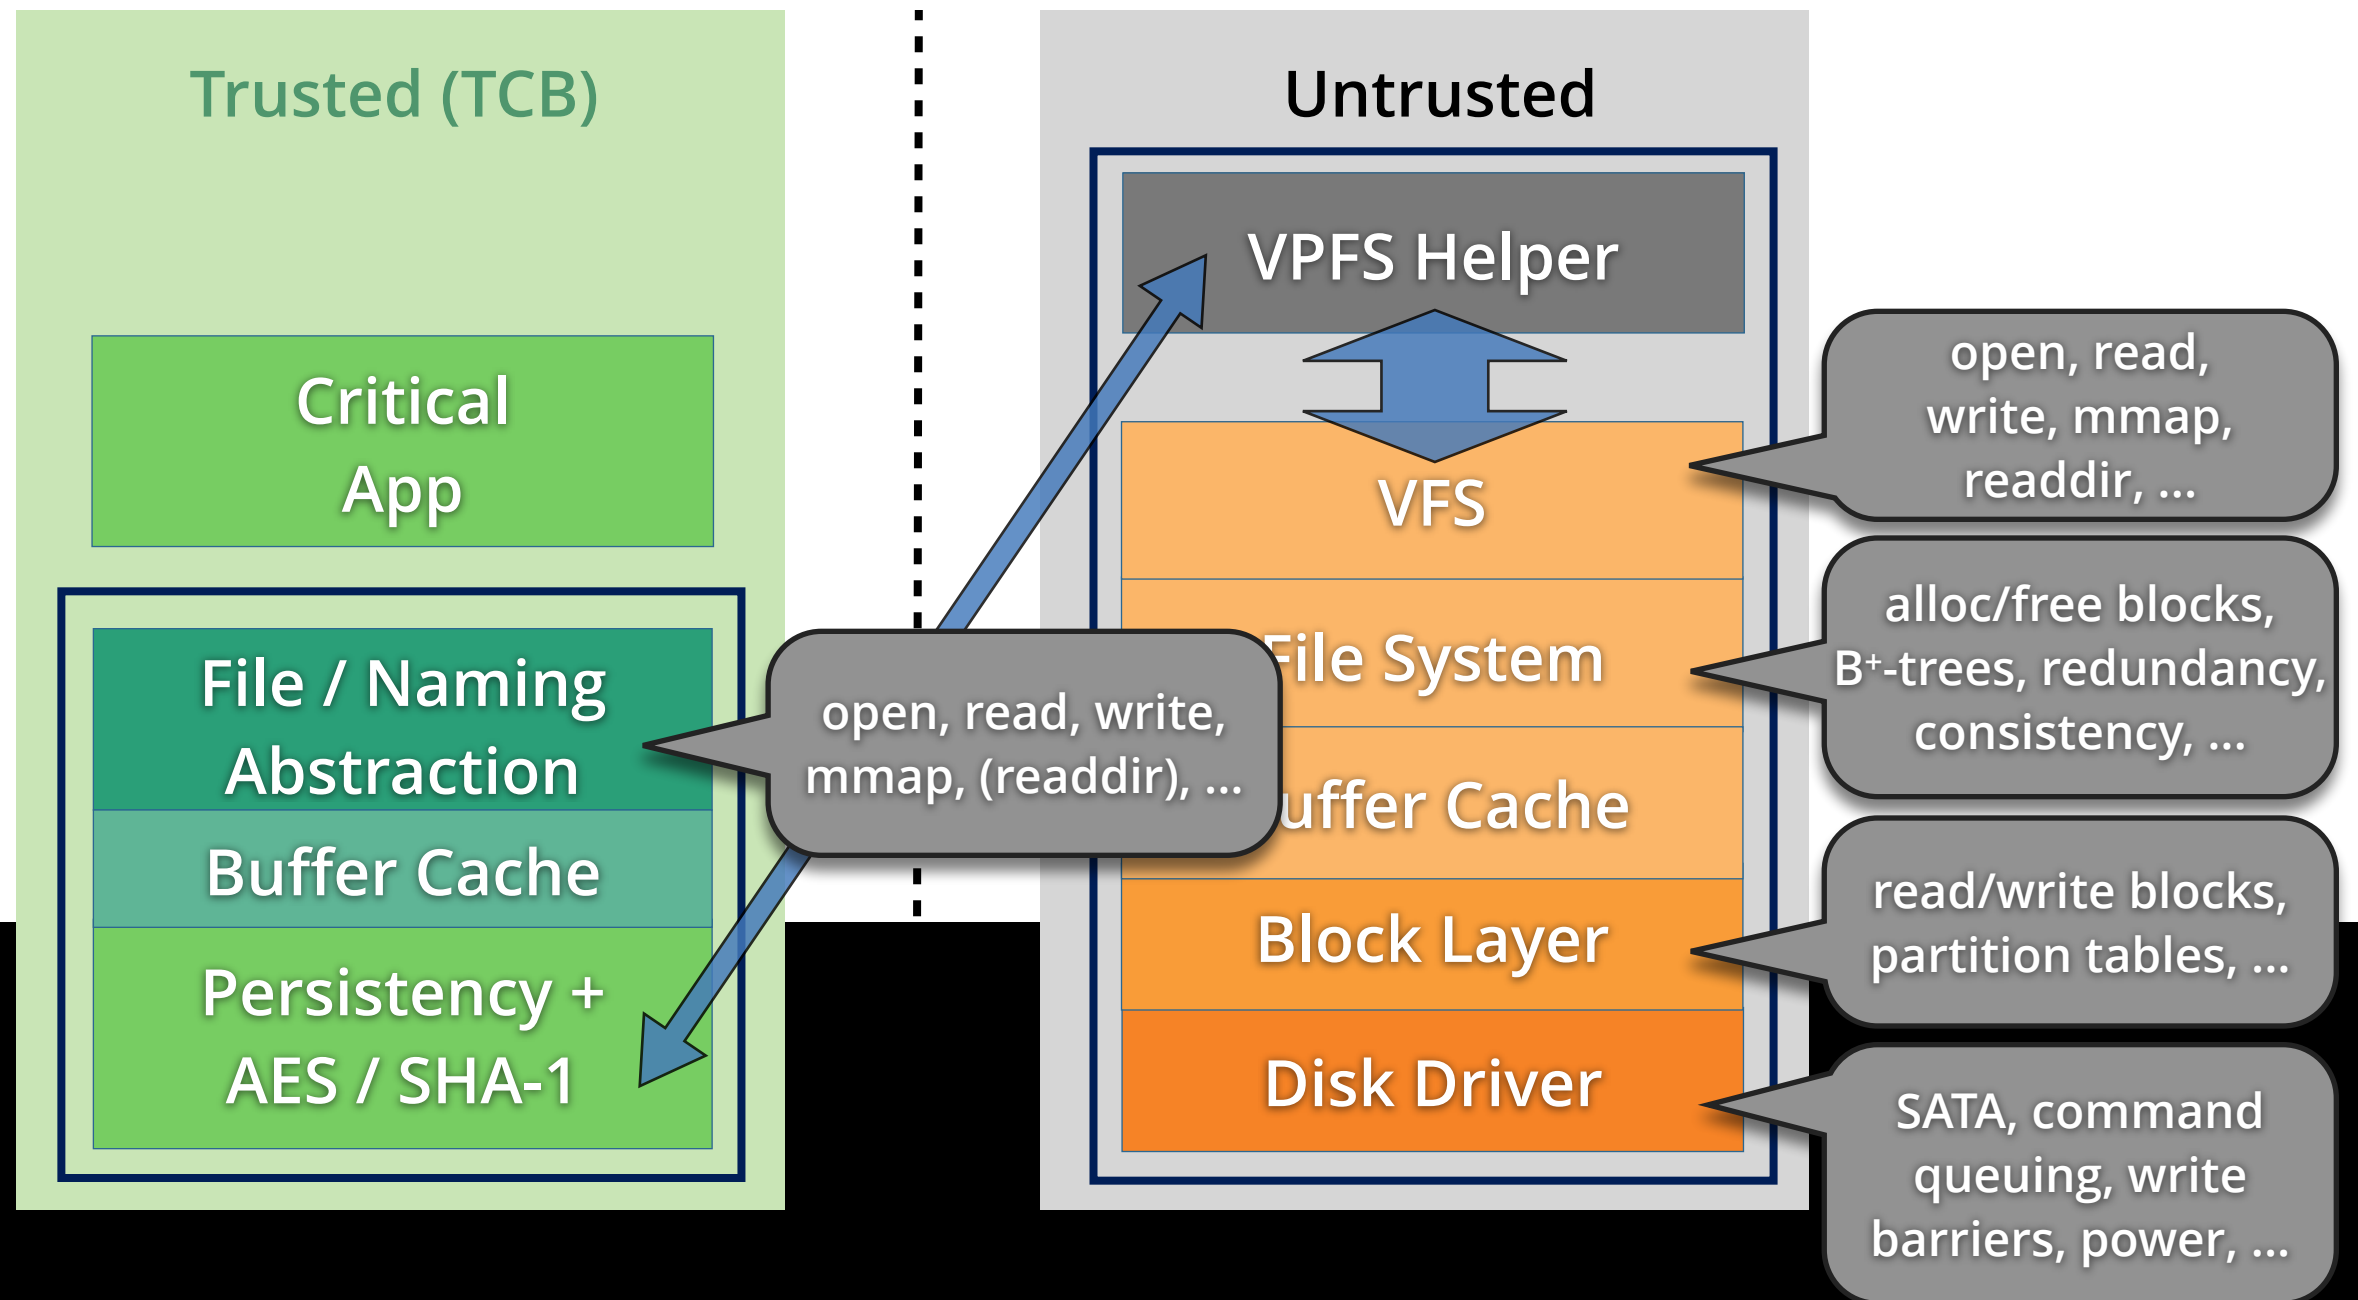


- Second end of design space:  
Protect individual files
  - Stacked file system
  - Encrypt all data and some metadata (directories, ...)
  - More flexibility for integrity






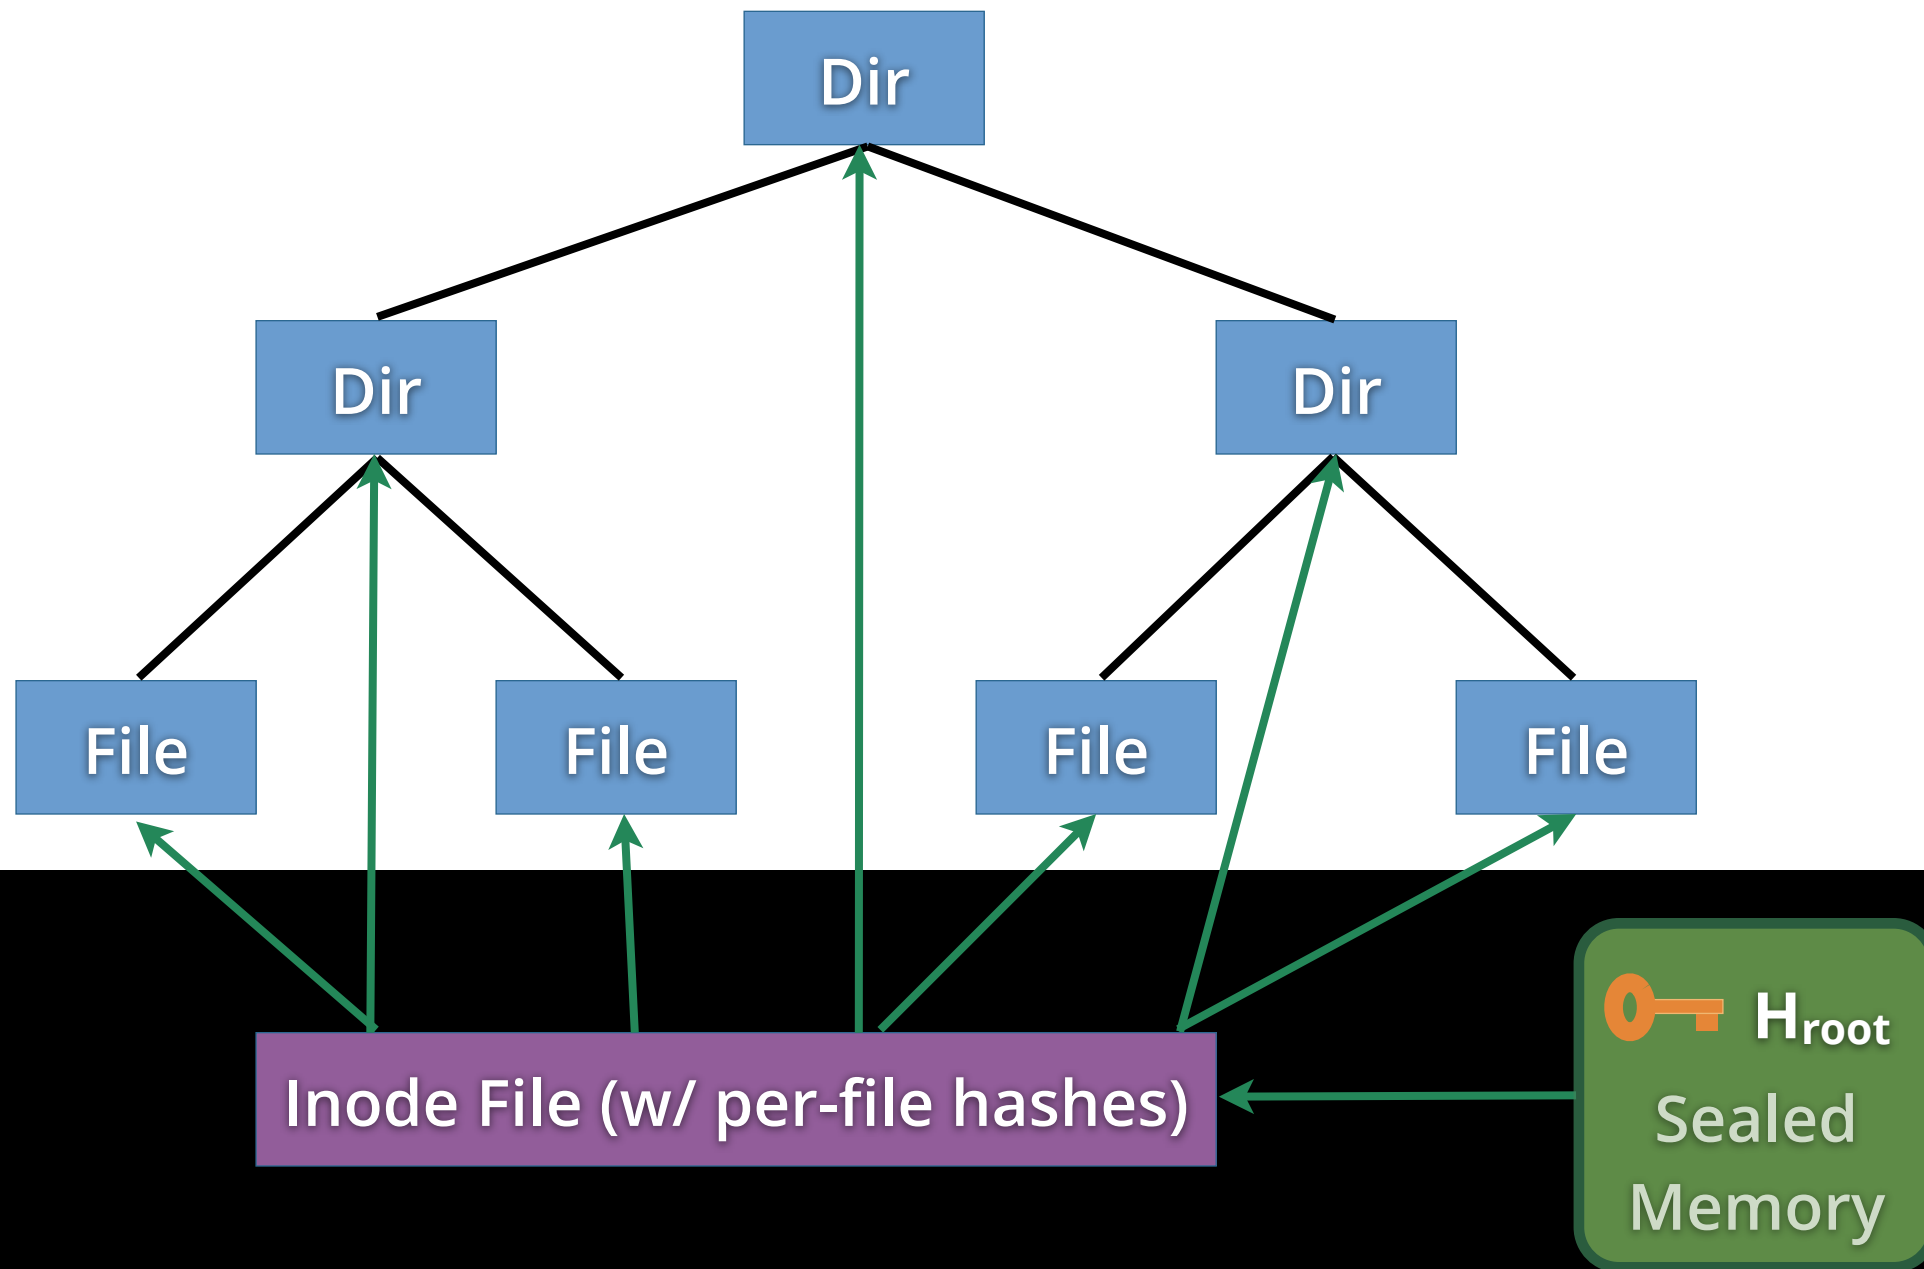
# TRUSTED WRAPPER



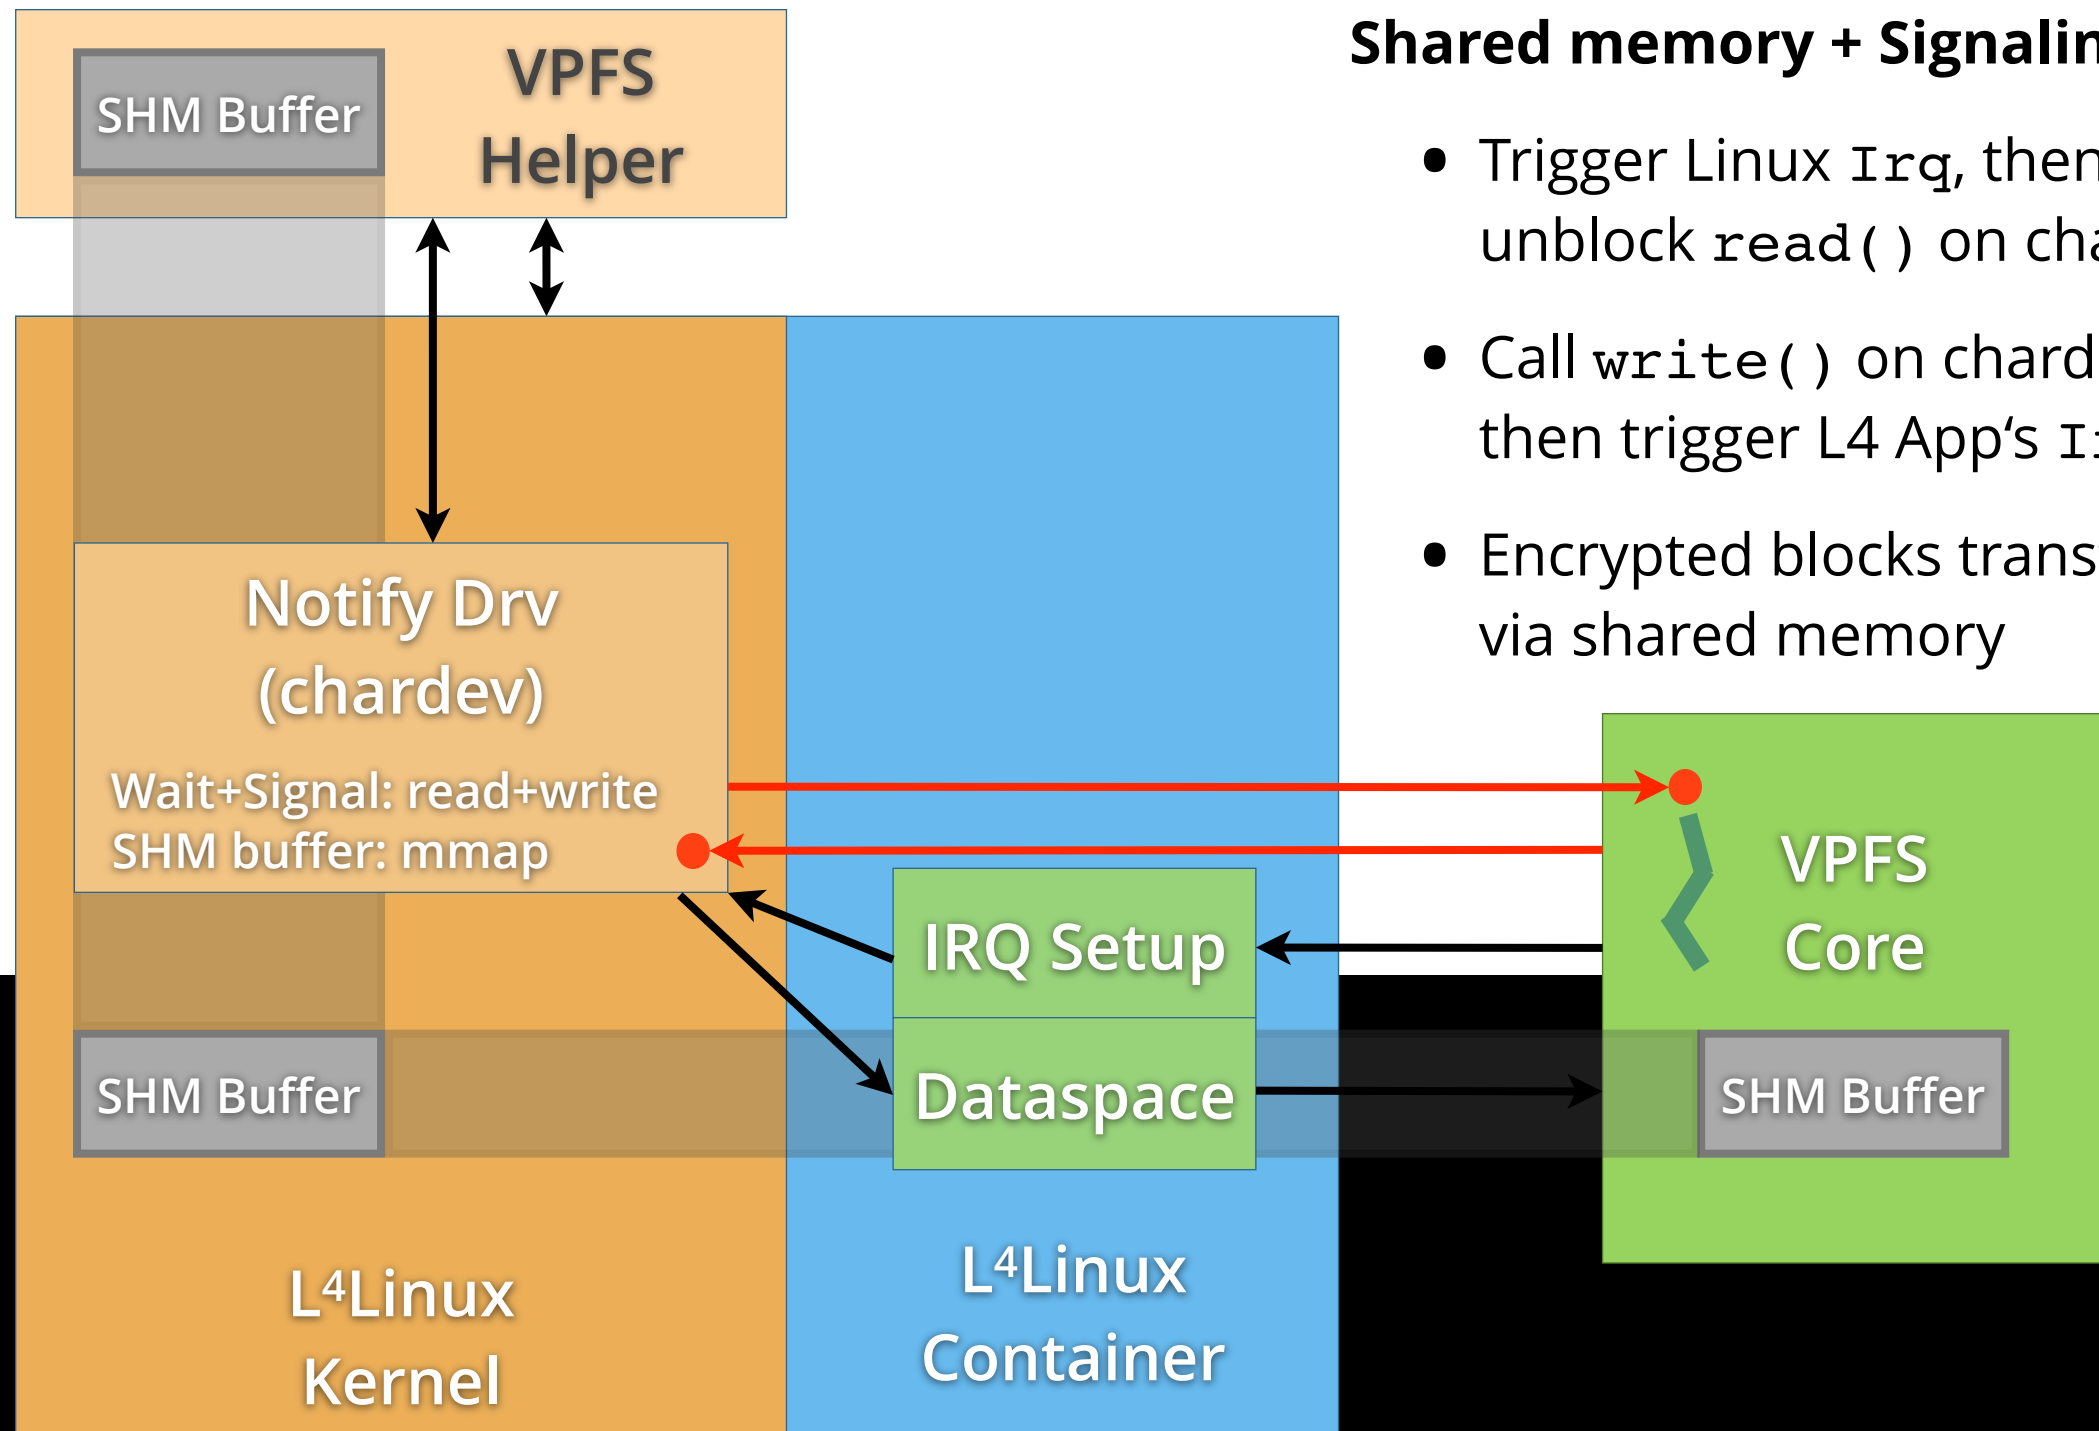
# VPFS APPROACH



- Trusted part of VPFS enforces security:
  - Encryption / decryption on the fly
  - Plaintext only in trusted buffer cache
  - Files in untrusted commodity file system store encrypted blocks



- VPFS reuses Linux file system stack:
  - Drivers, block device layer
  - Optimizations (buffer cache, read ahead, write batching, ...)
  - Allocate / free disk storage for files




## Shared memory + Signaling:

- Trigger Linux Irq, then unblock `read()` on chardev
- Call `write()` on chardev, then trigger L4 App's Irq
- Encrypted blocks transferred via shared memory

- Trusted wrappers for file systems work!
- VPFS is general purpose file system
- Significant reduction in code size:
  - Untrusted Linux file system stack comprises **50,000+** SLOC
  - VPFS adds **4 000** to **4 600** SLOC to

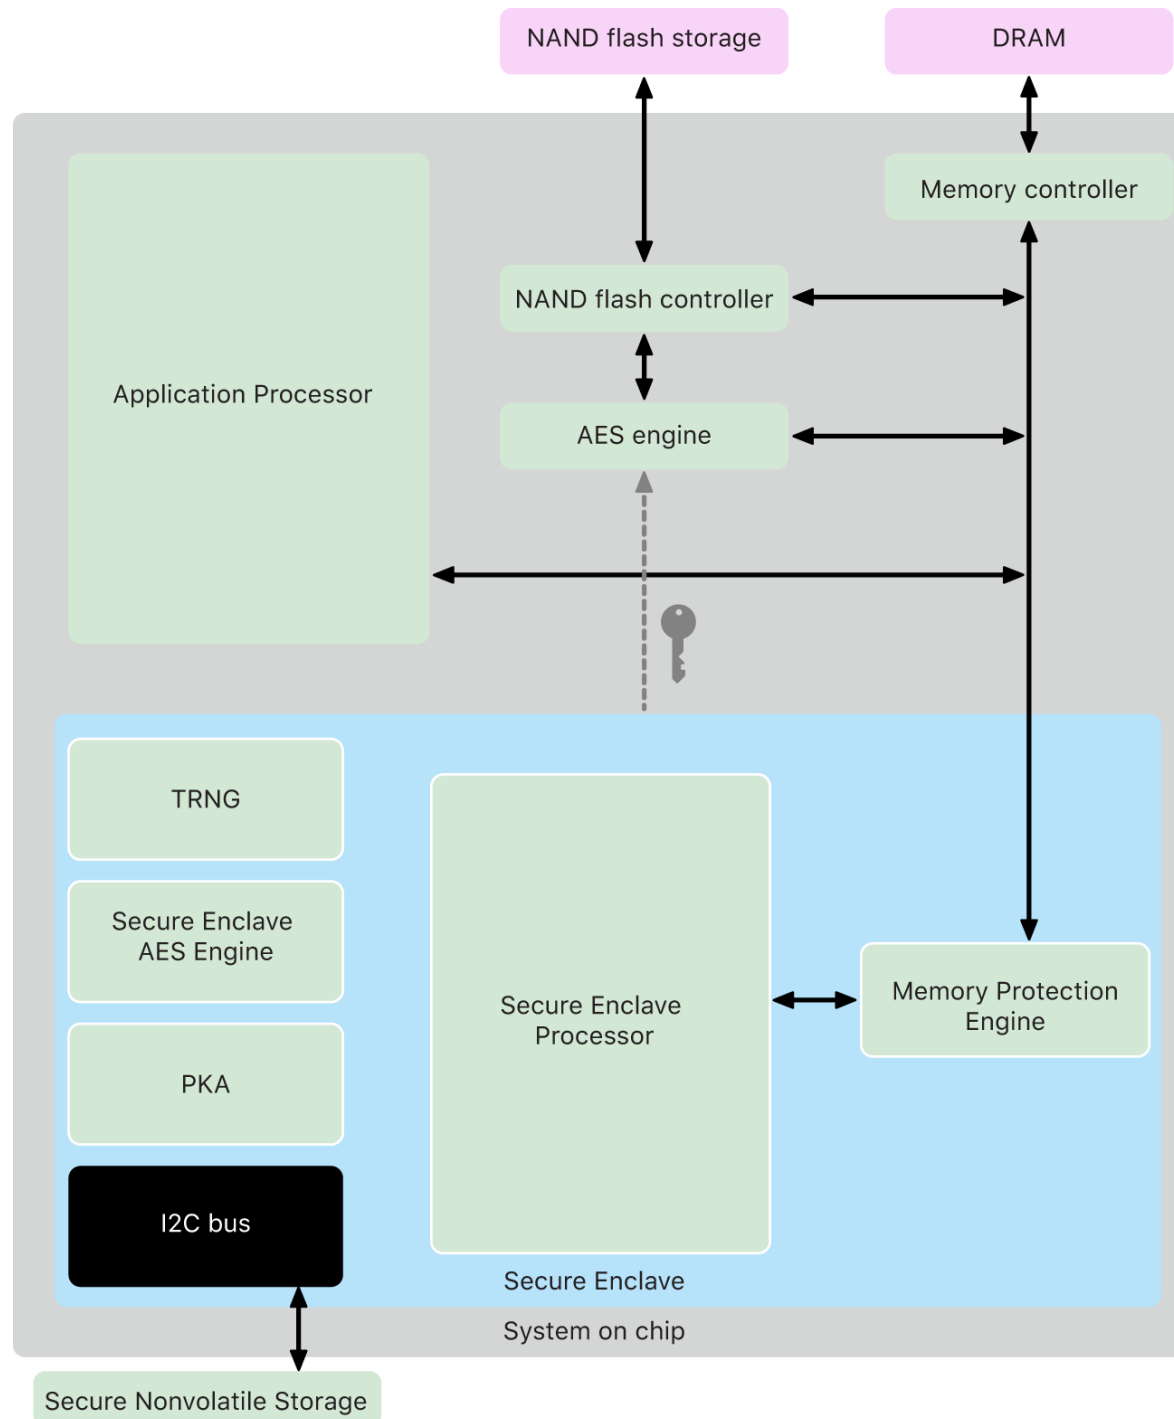
- Secure reuse of untrusted legacy infrastructure
- Split apps + OS services for smaller TCB
- Nizza secure system architecture:
  - Strong isolation





# **BRIEFLY: HARDWARE ISOLATION**

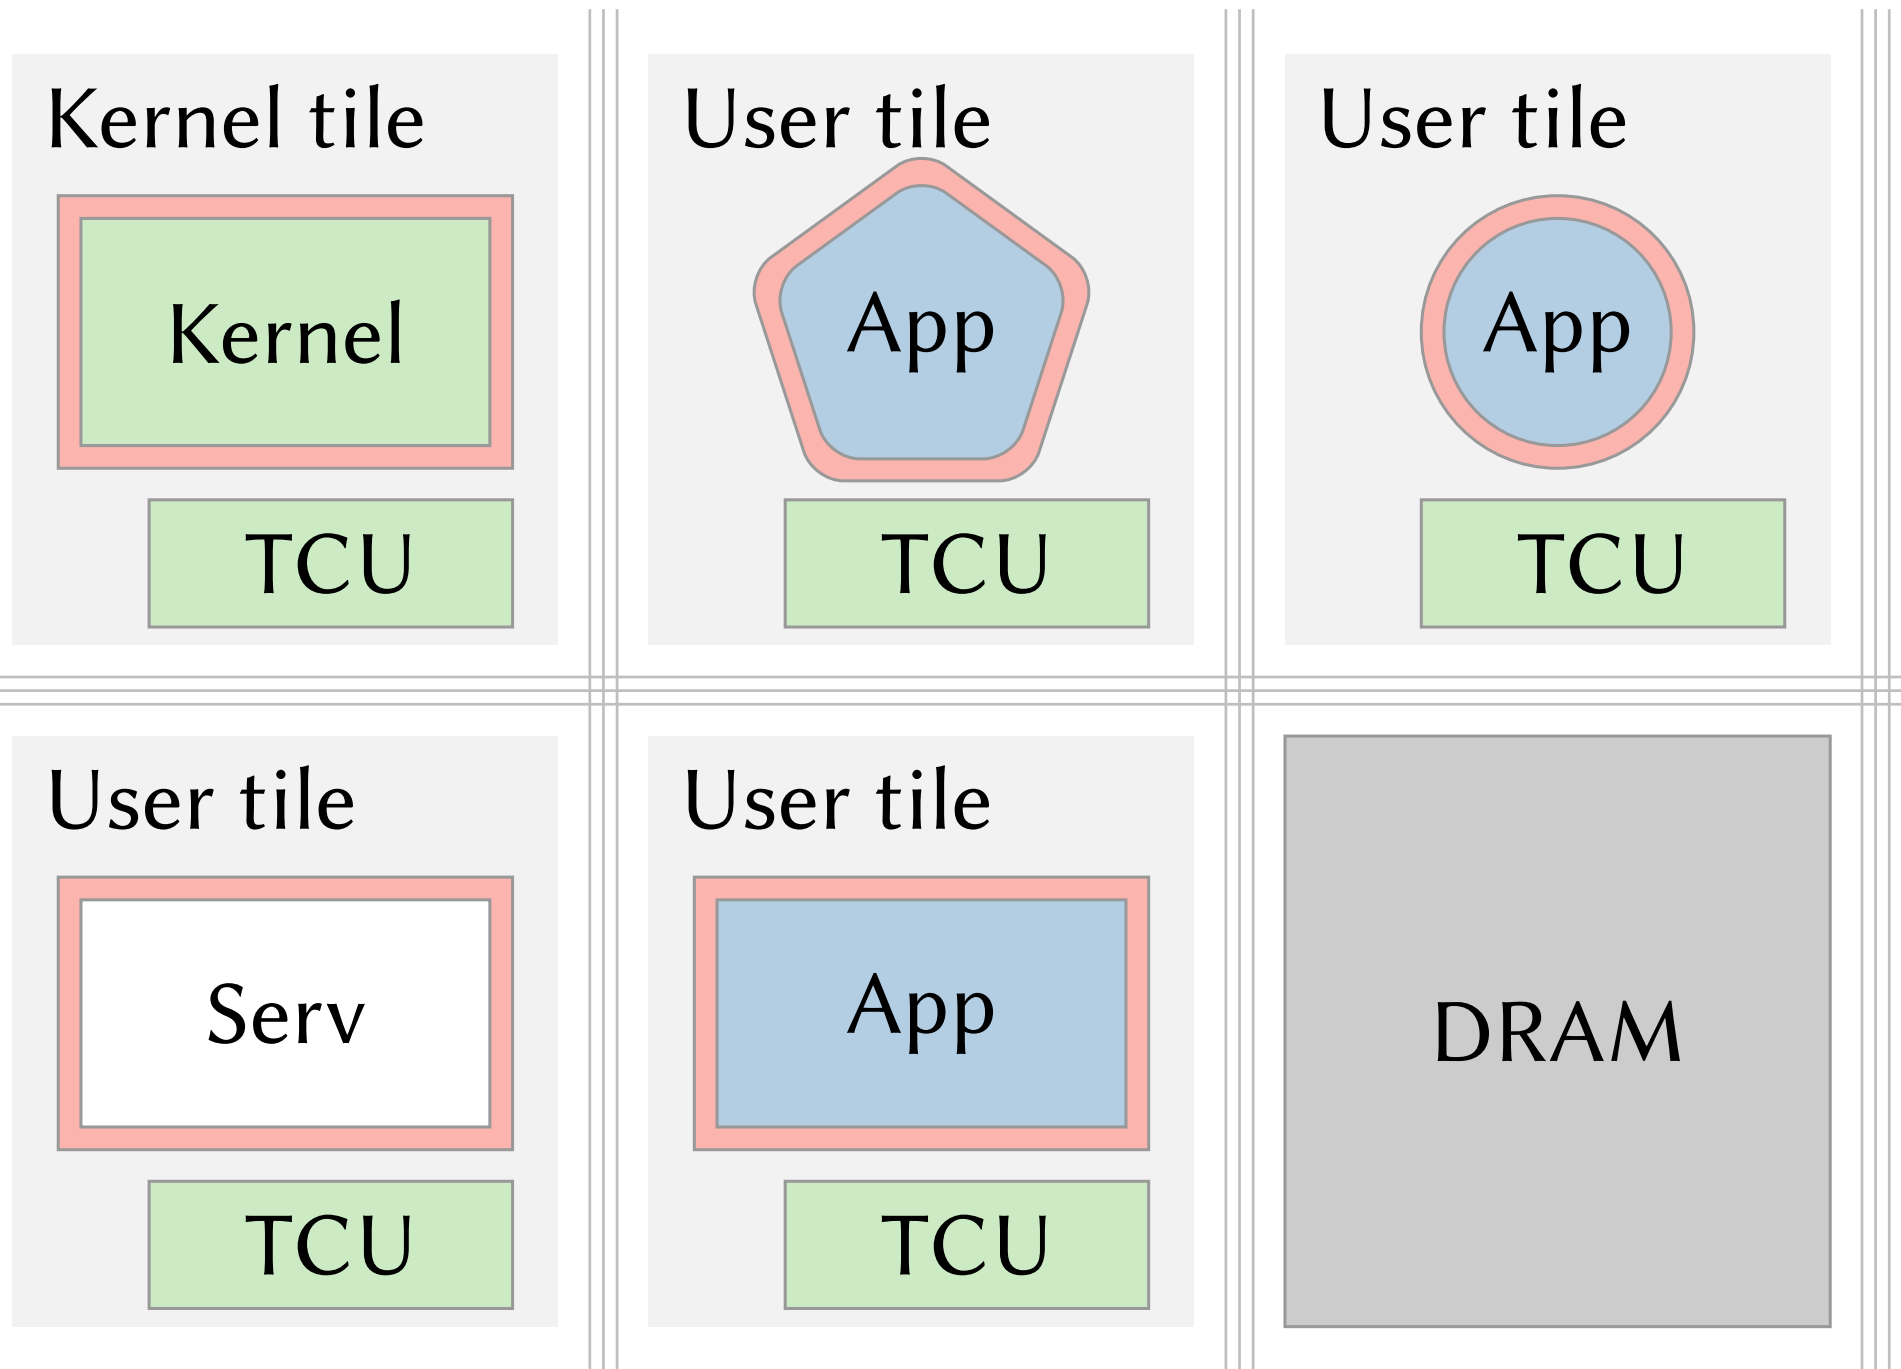
# APPLE SECURE ENCLAVE

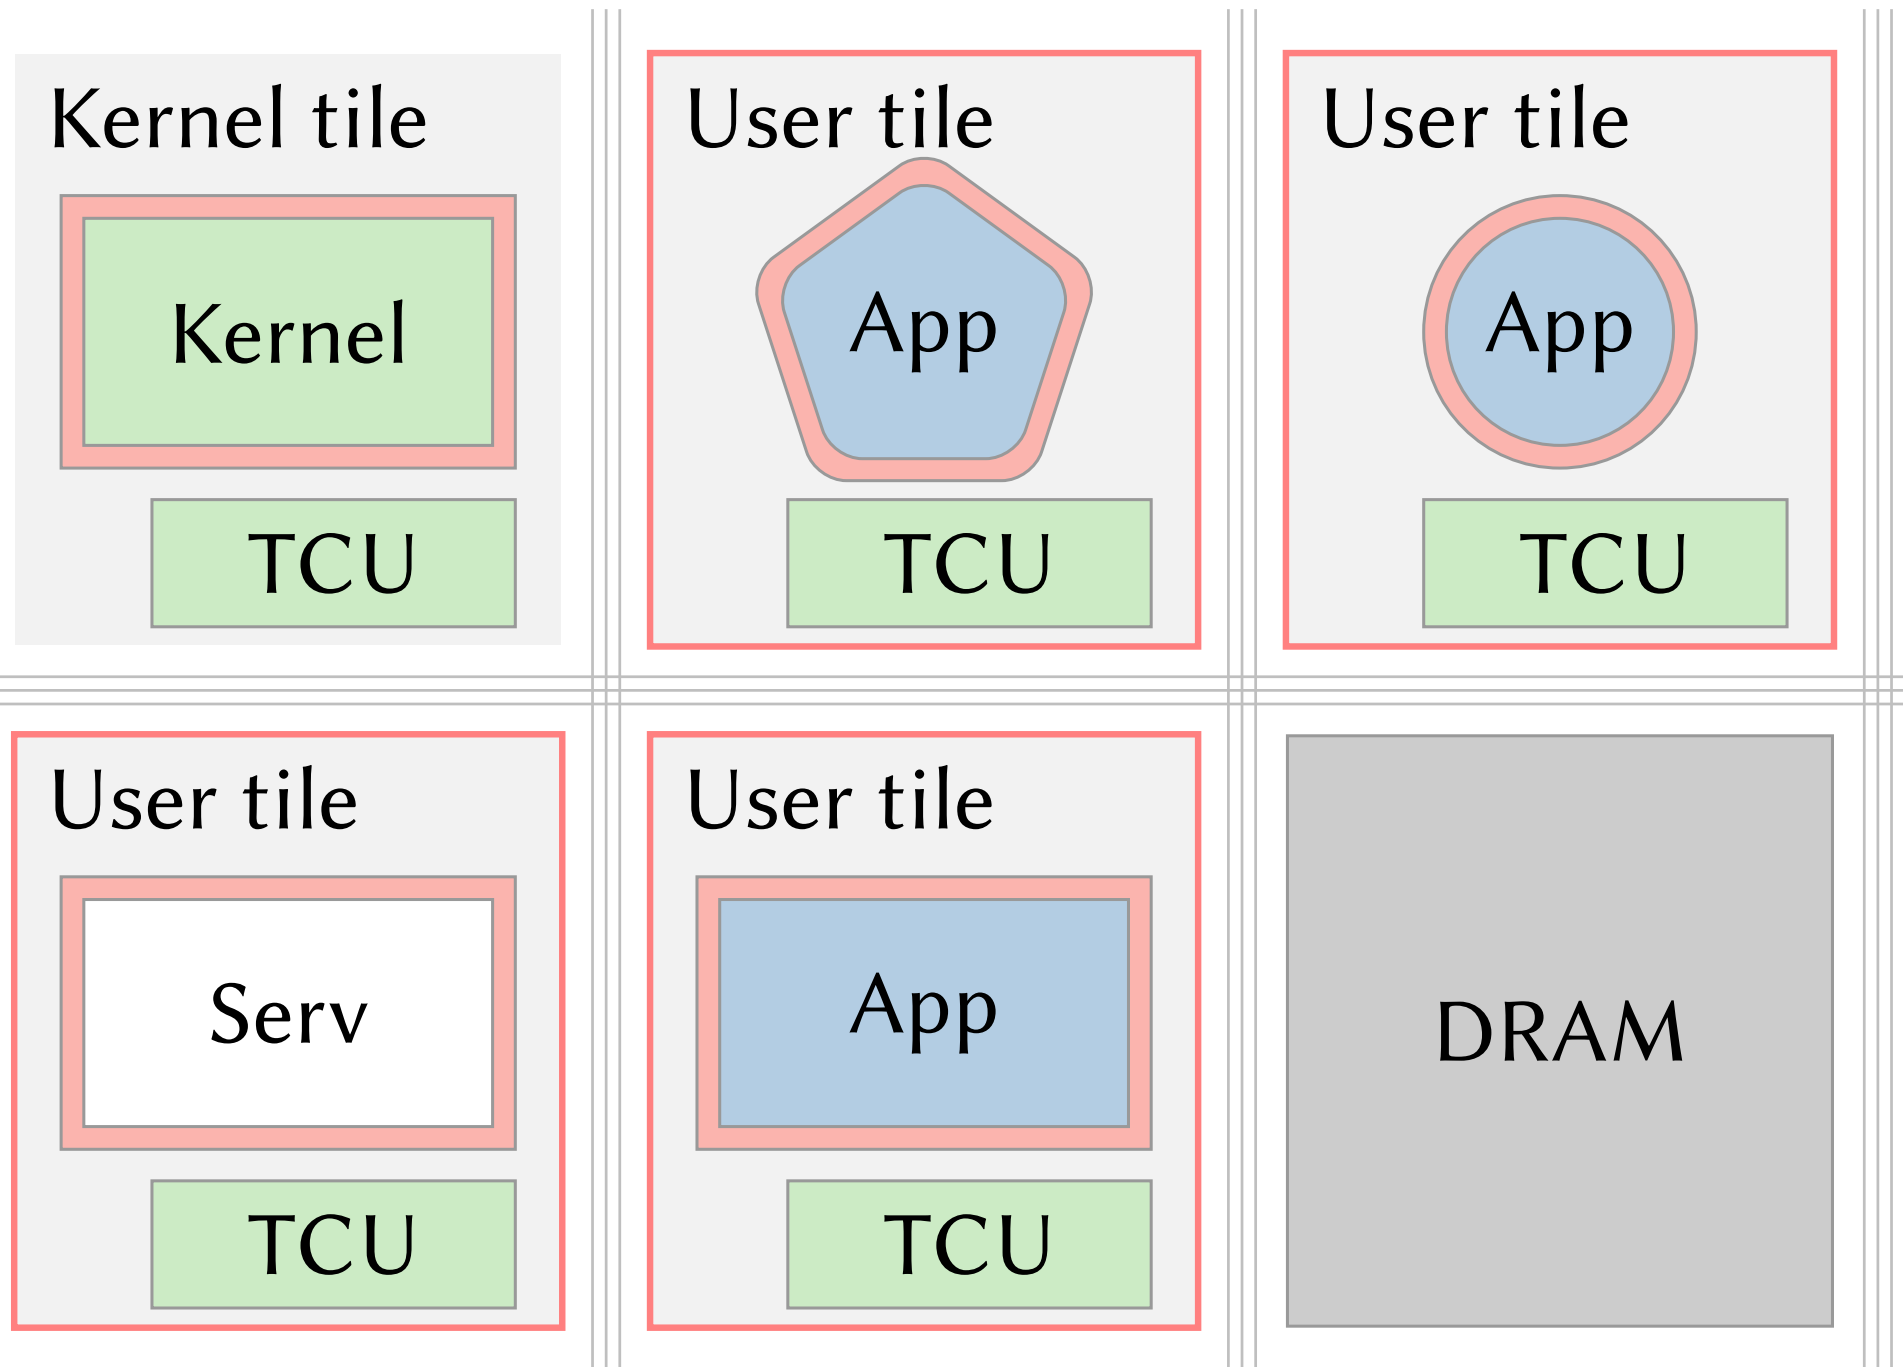


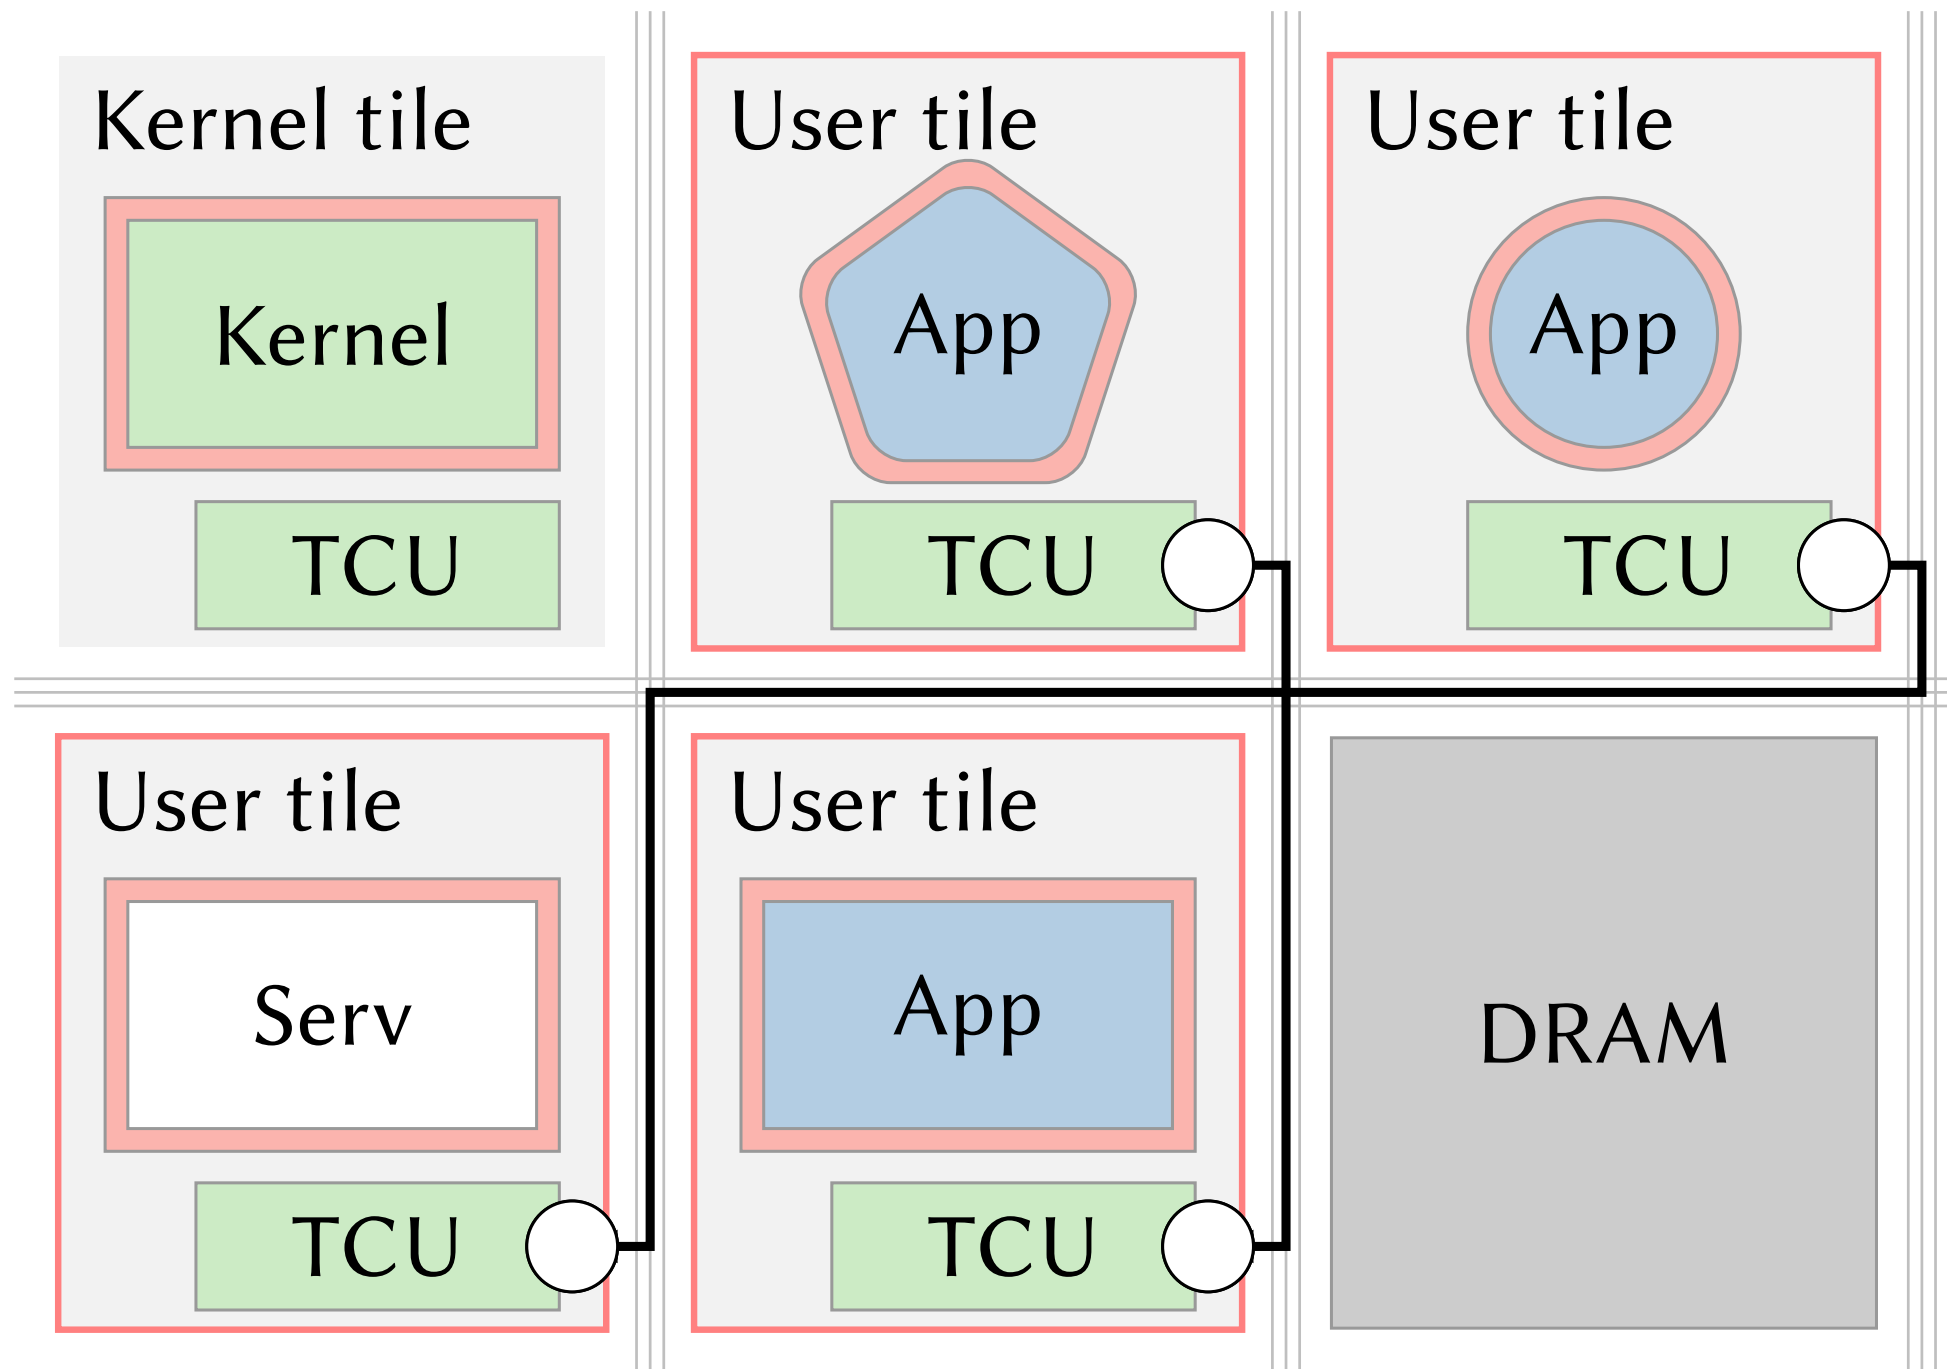
Apple devices have "Secure Enclave Processor (SEP)" running a dedicated service OS fully isolated from from the application processor hardware.

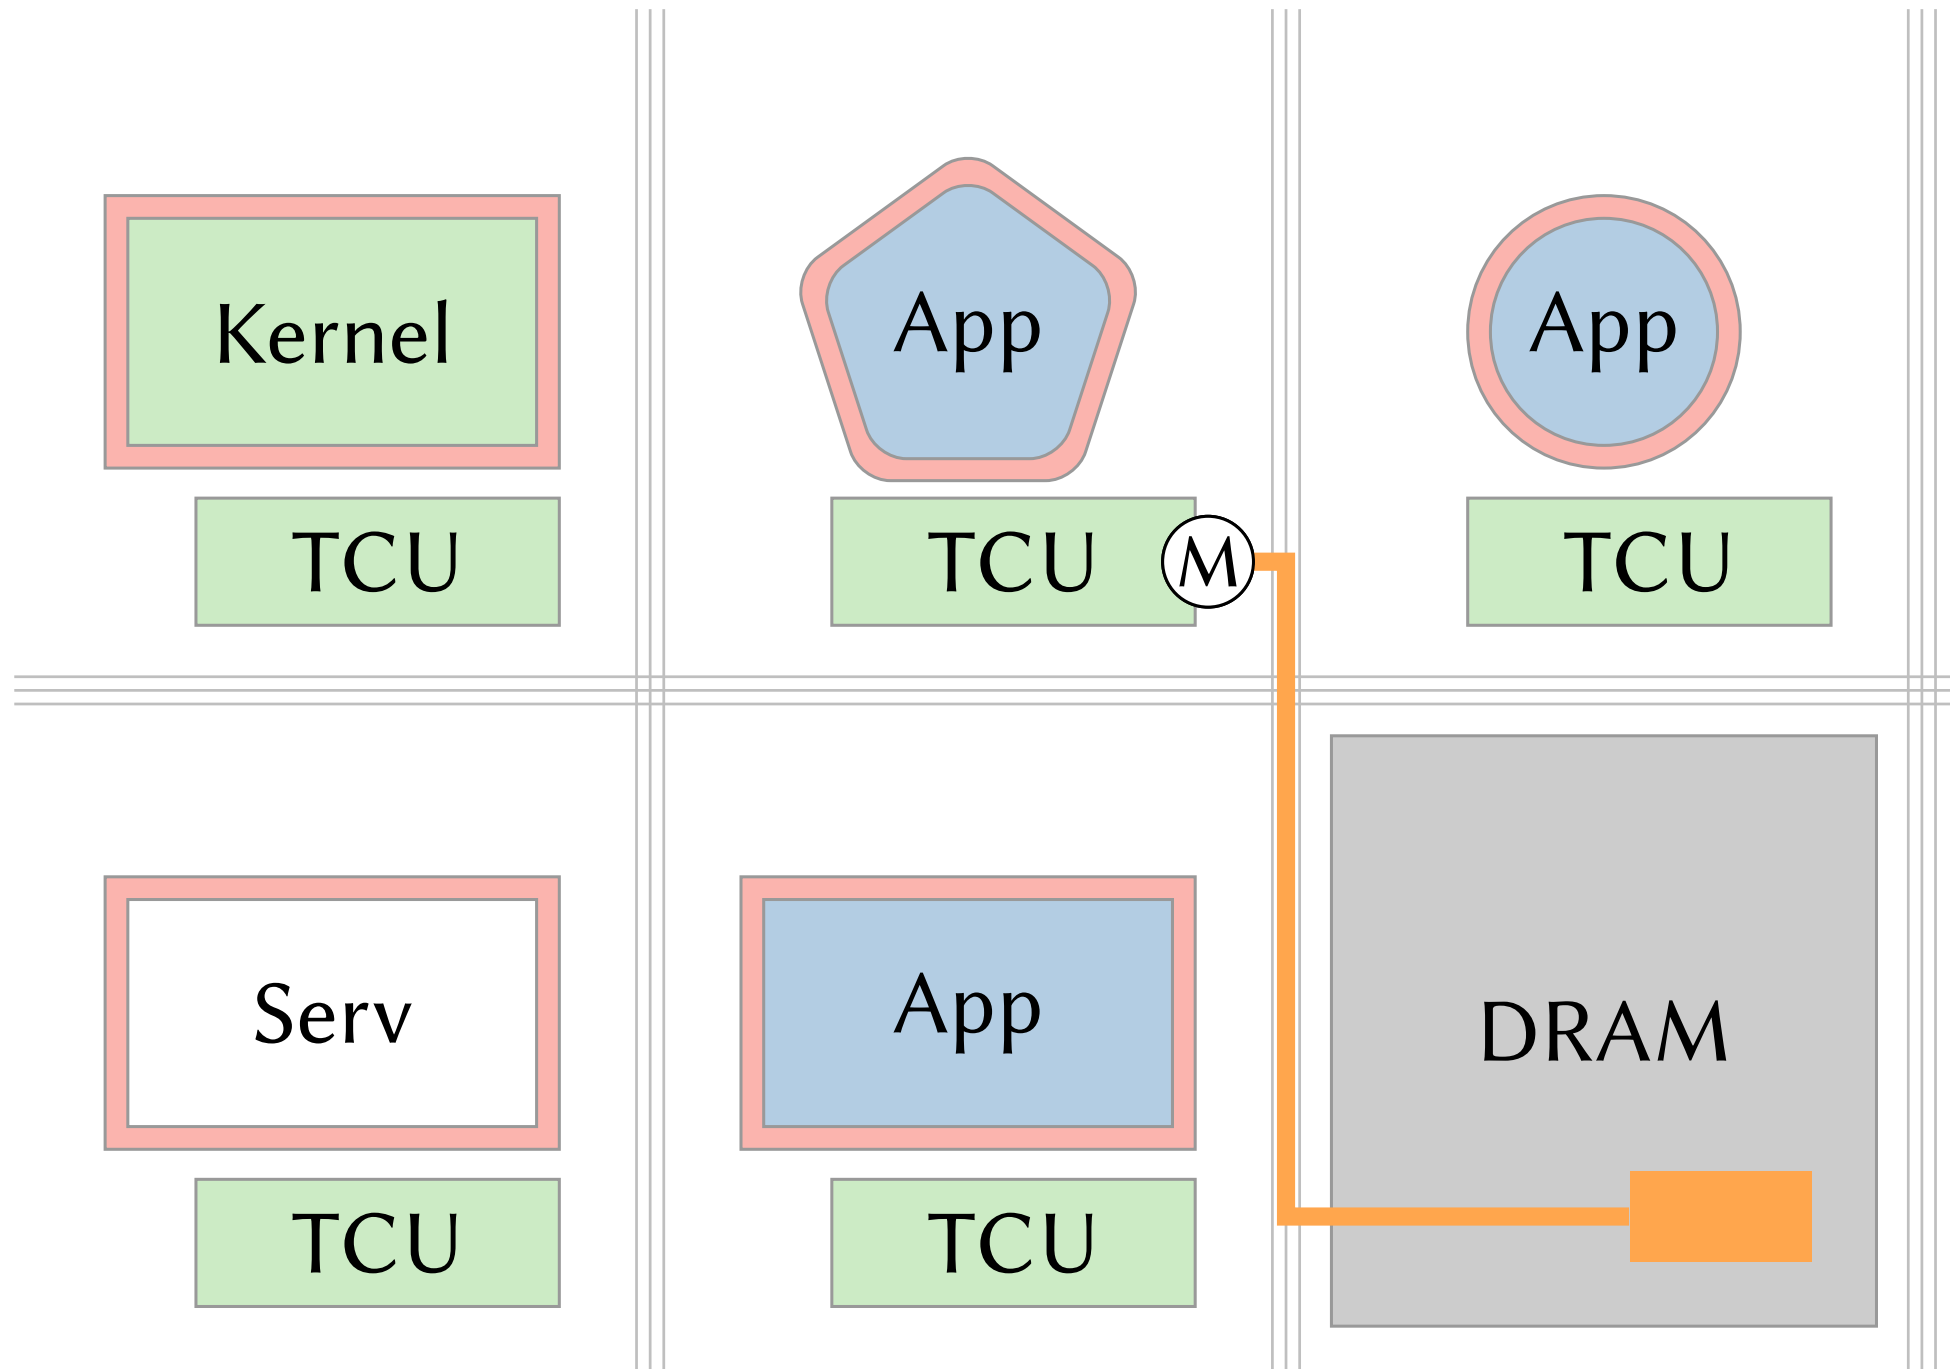
**The SEP runs sepOS:**

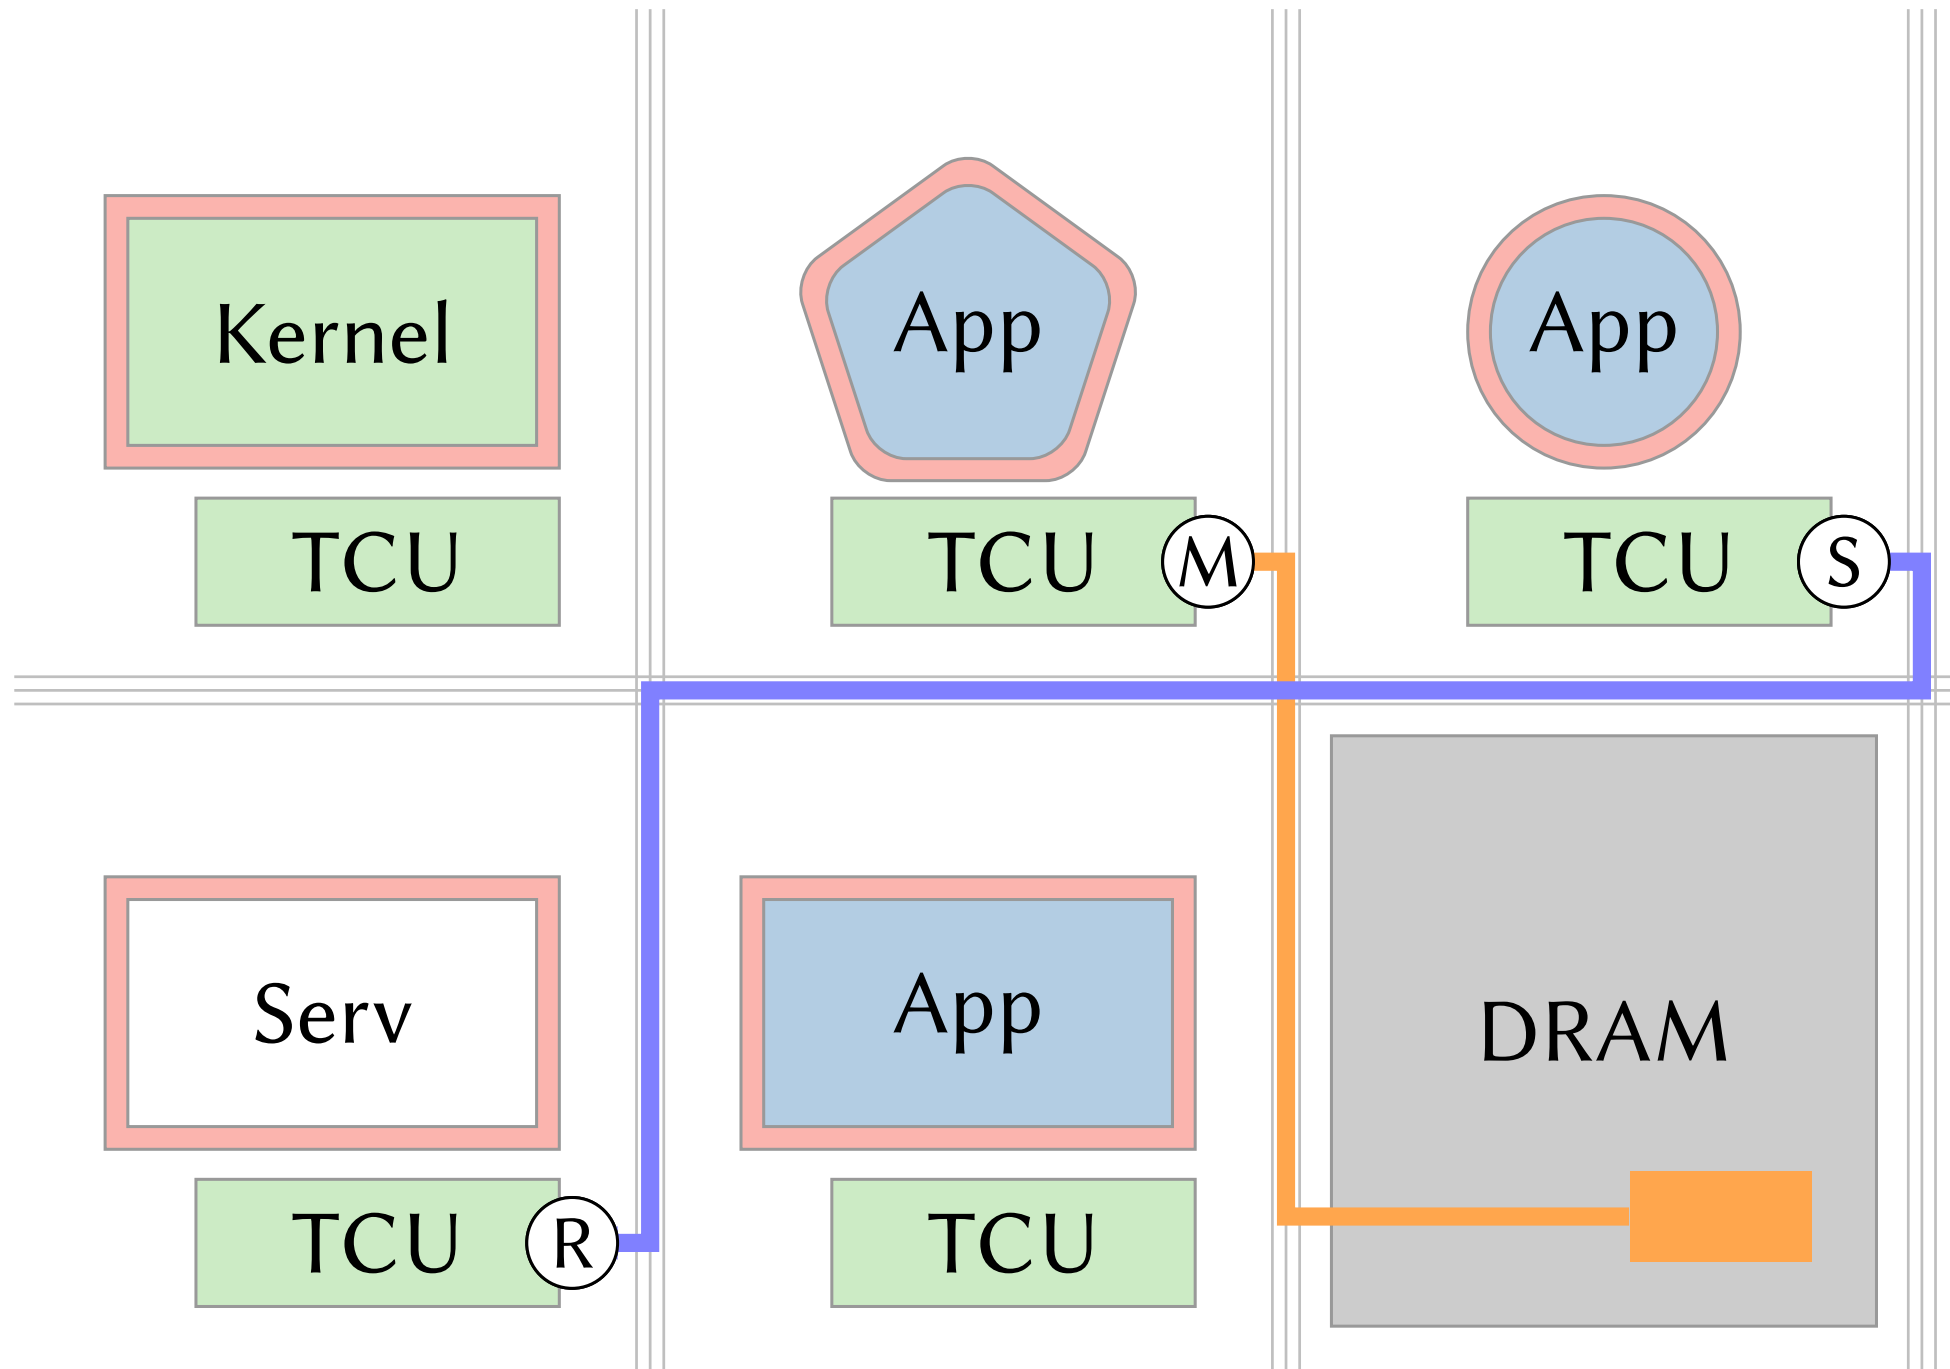
Source: Apple Support Documentation  
<https://support.apple.com/guide/security/secure-enclave-sec59b0b31ff/web>











- [1] <http://www.heise.de/newsticker/Month-of-Kernel-Bugs-Ein-Zwischenstand--/meldung/81454>
- [2] <http://projects.info-pull.com/mokb/>
- [3] Carsten Weinhold and Hermann Härtig, „VPFS: Building a Virtual Private File System with a Small Trusted Computing Base“, Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems, April 2008, Glasgow, Scotland UK
- [4] Carsten Weinhold and Hermann Härtig, „jVPFS: Adding Robustness to a Secure Stacked File System with Untrusted Local Storage Components“, Proceedings of the 2011 USENIX Annual Technical Conference, Portland, OR, USA, June 2011
- [5] Norman Feske and Christian Helmuth, „A Nitpicker's guide to a minimal-complexity secure GUI“, ACSAC '05: Proceedings of the 21st Annual Computer Security Applications Conference, 2005, Washington, DC, USA
- [6] Christian Helmuth, Alexander Warg, Norman Feske, „Mikro-SINA - Hands-on Experiences with the Nizza Security Architecture“, D.A.CH Security 2005, 2005, Darmstadt, Germany