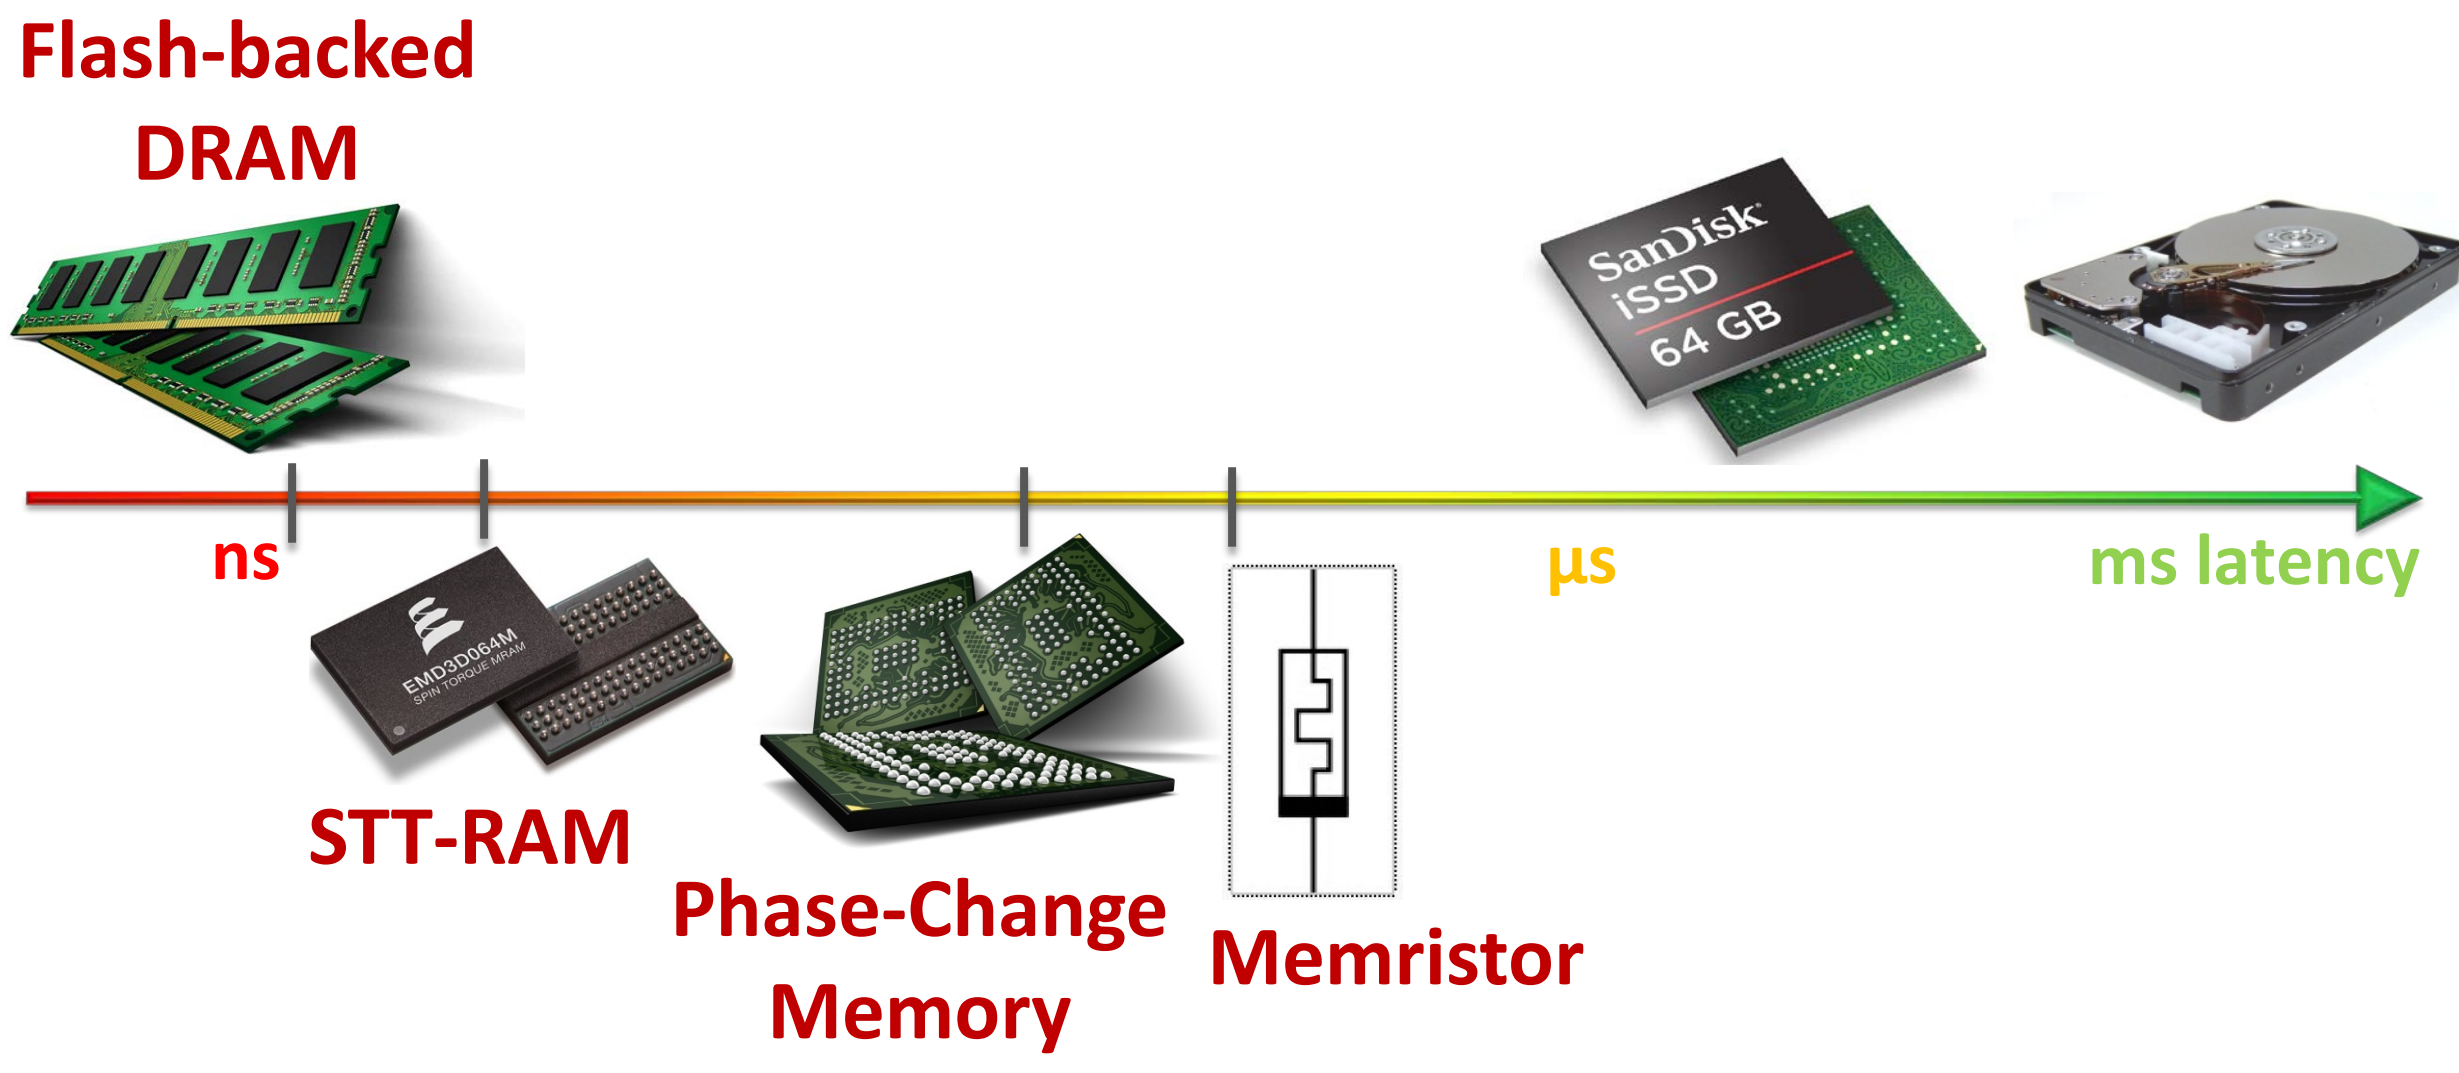


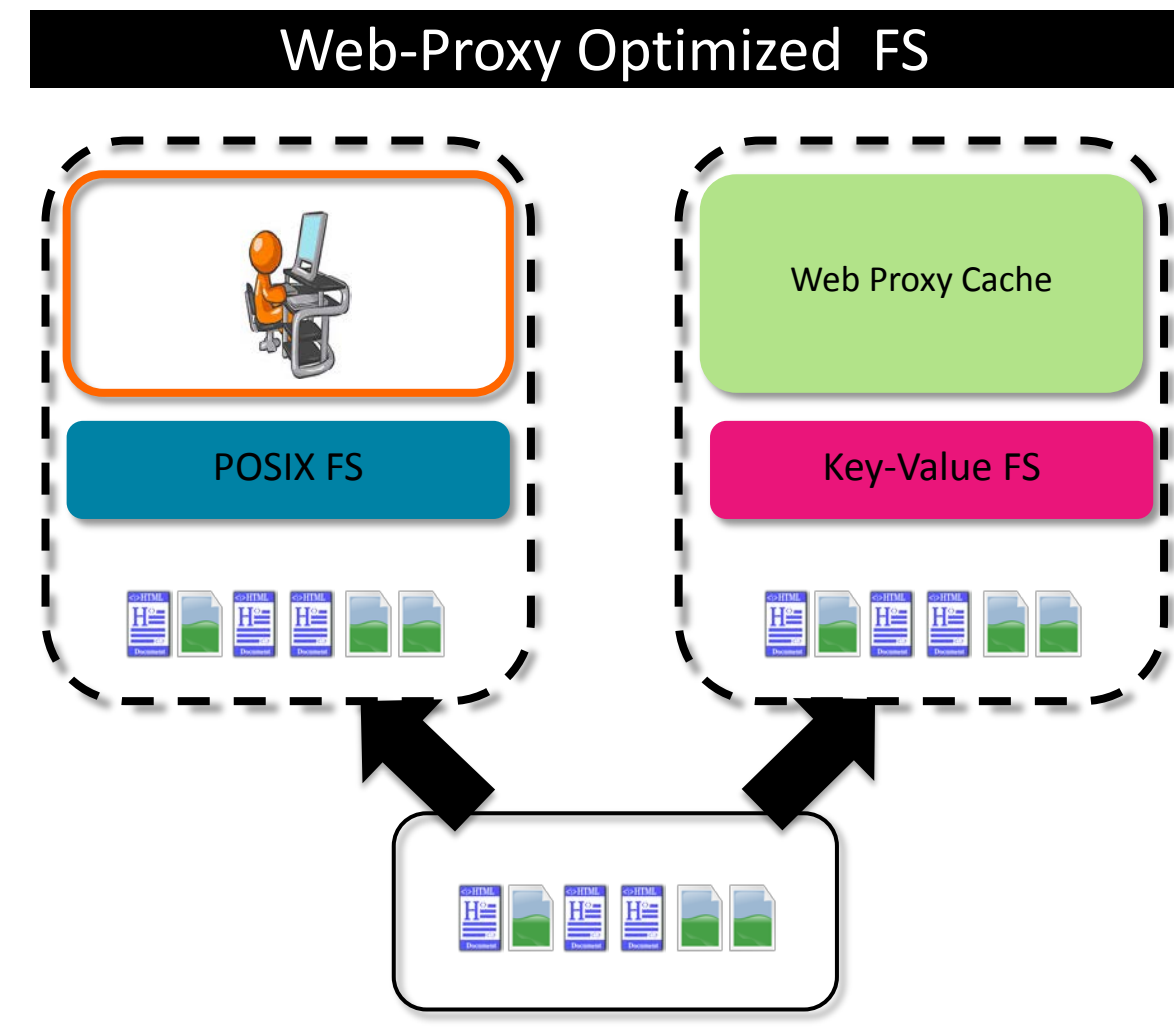
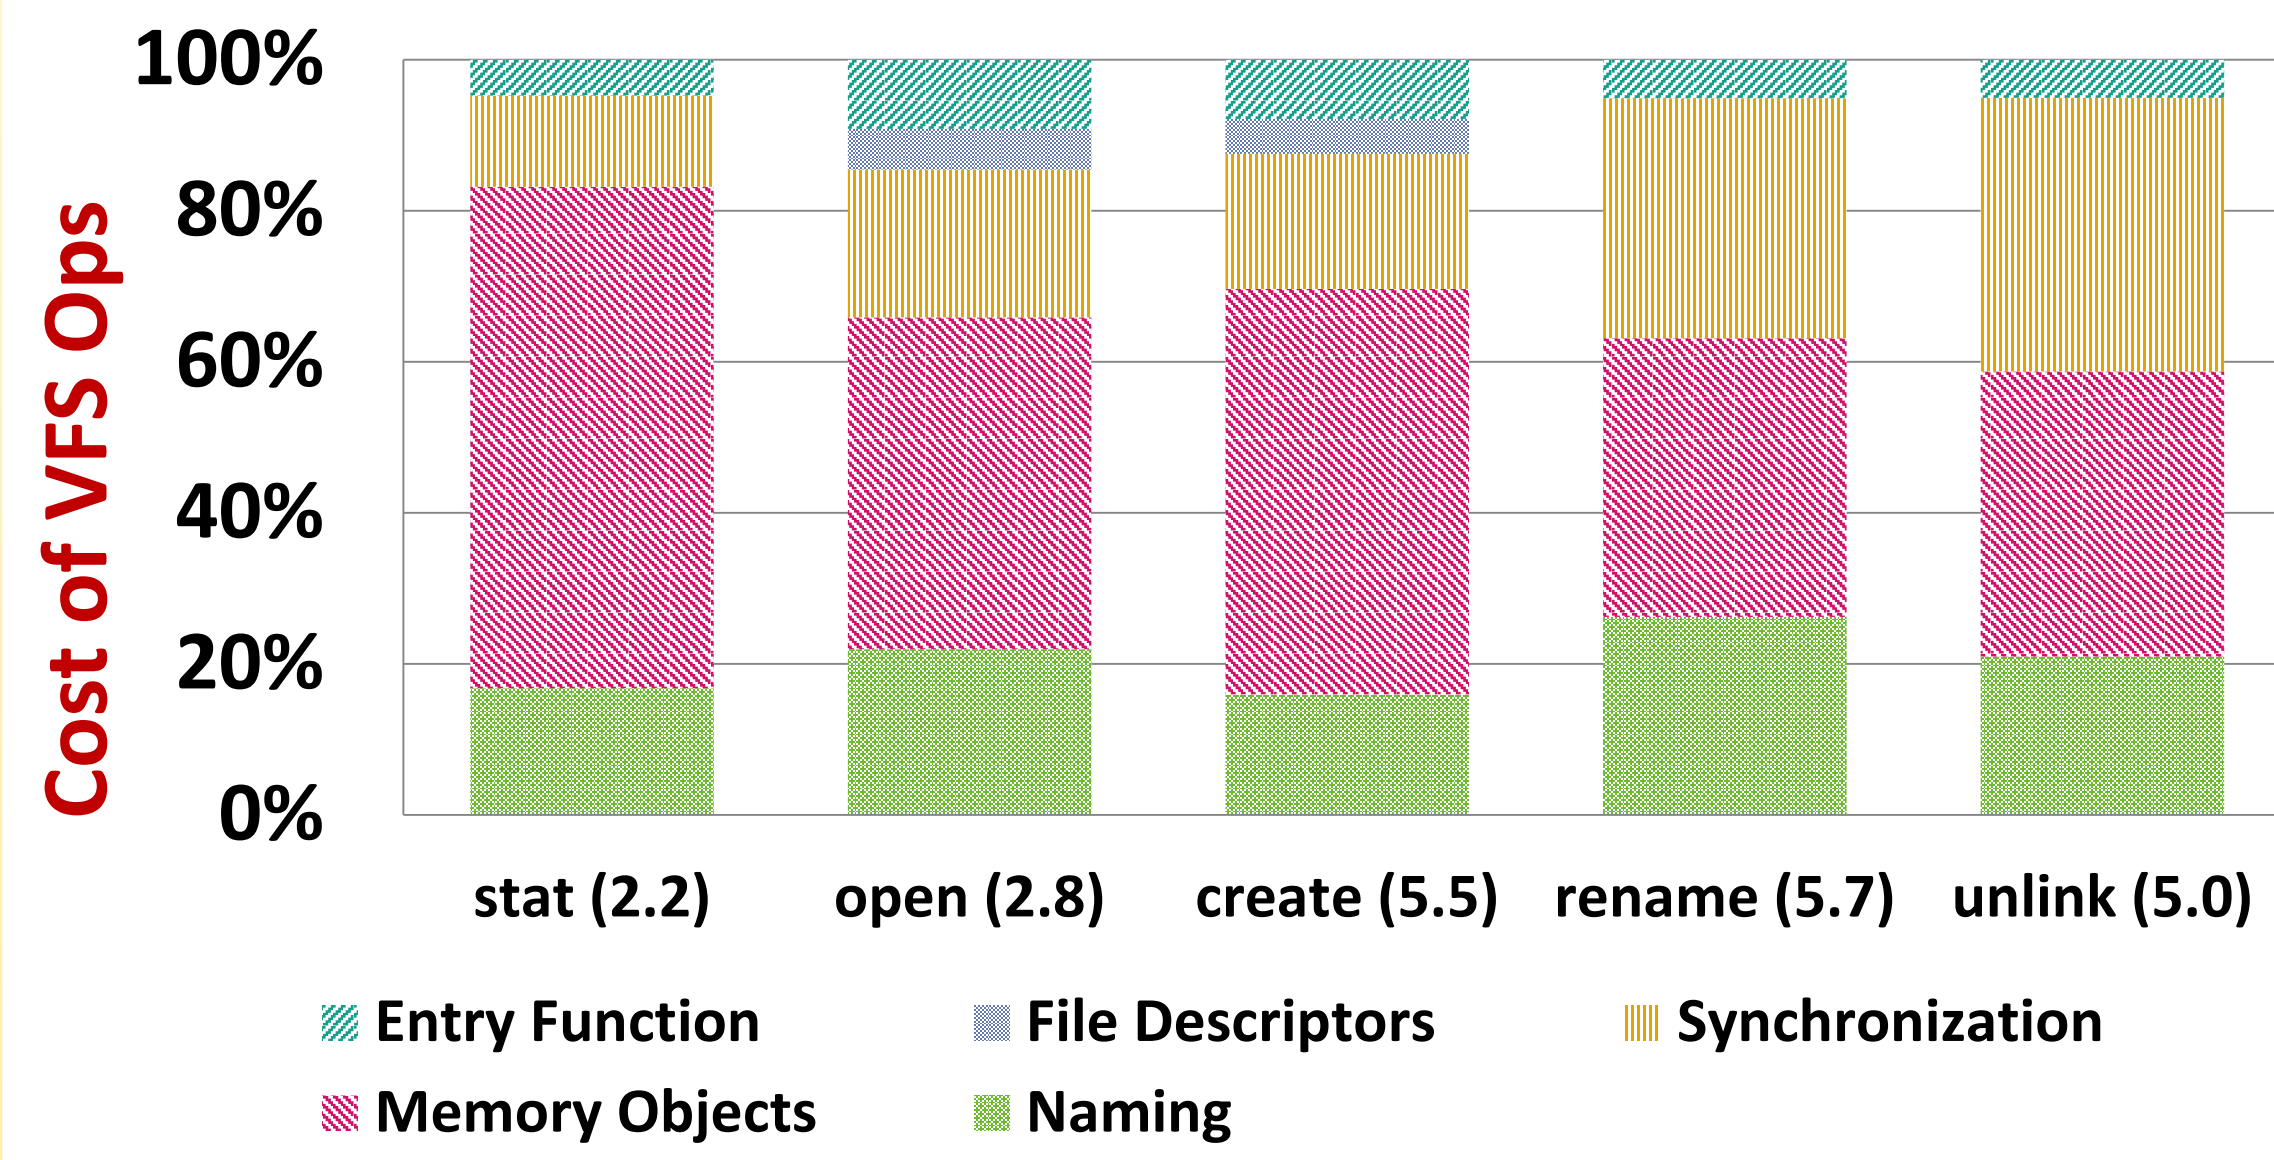
Aerie: Flexible File-System Interfaces to Storage-Class Memory (SCM)

STORAGE-CLASS MEMORY (SCM)



★ **Fast persistent memory !**

THE COST OF FILE ABSTRACTION



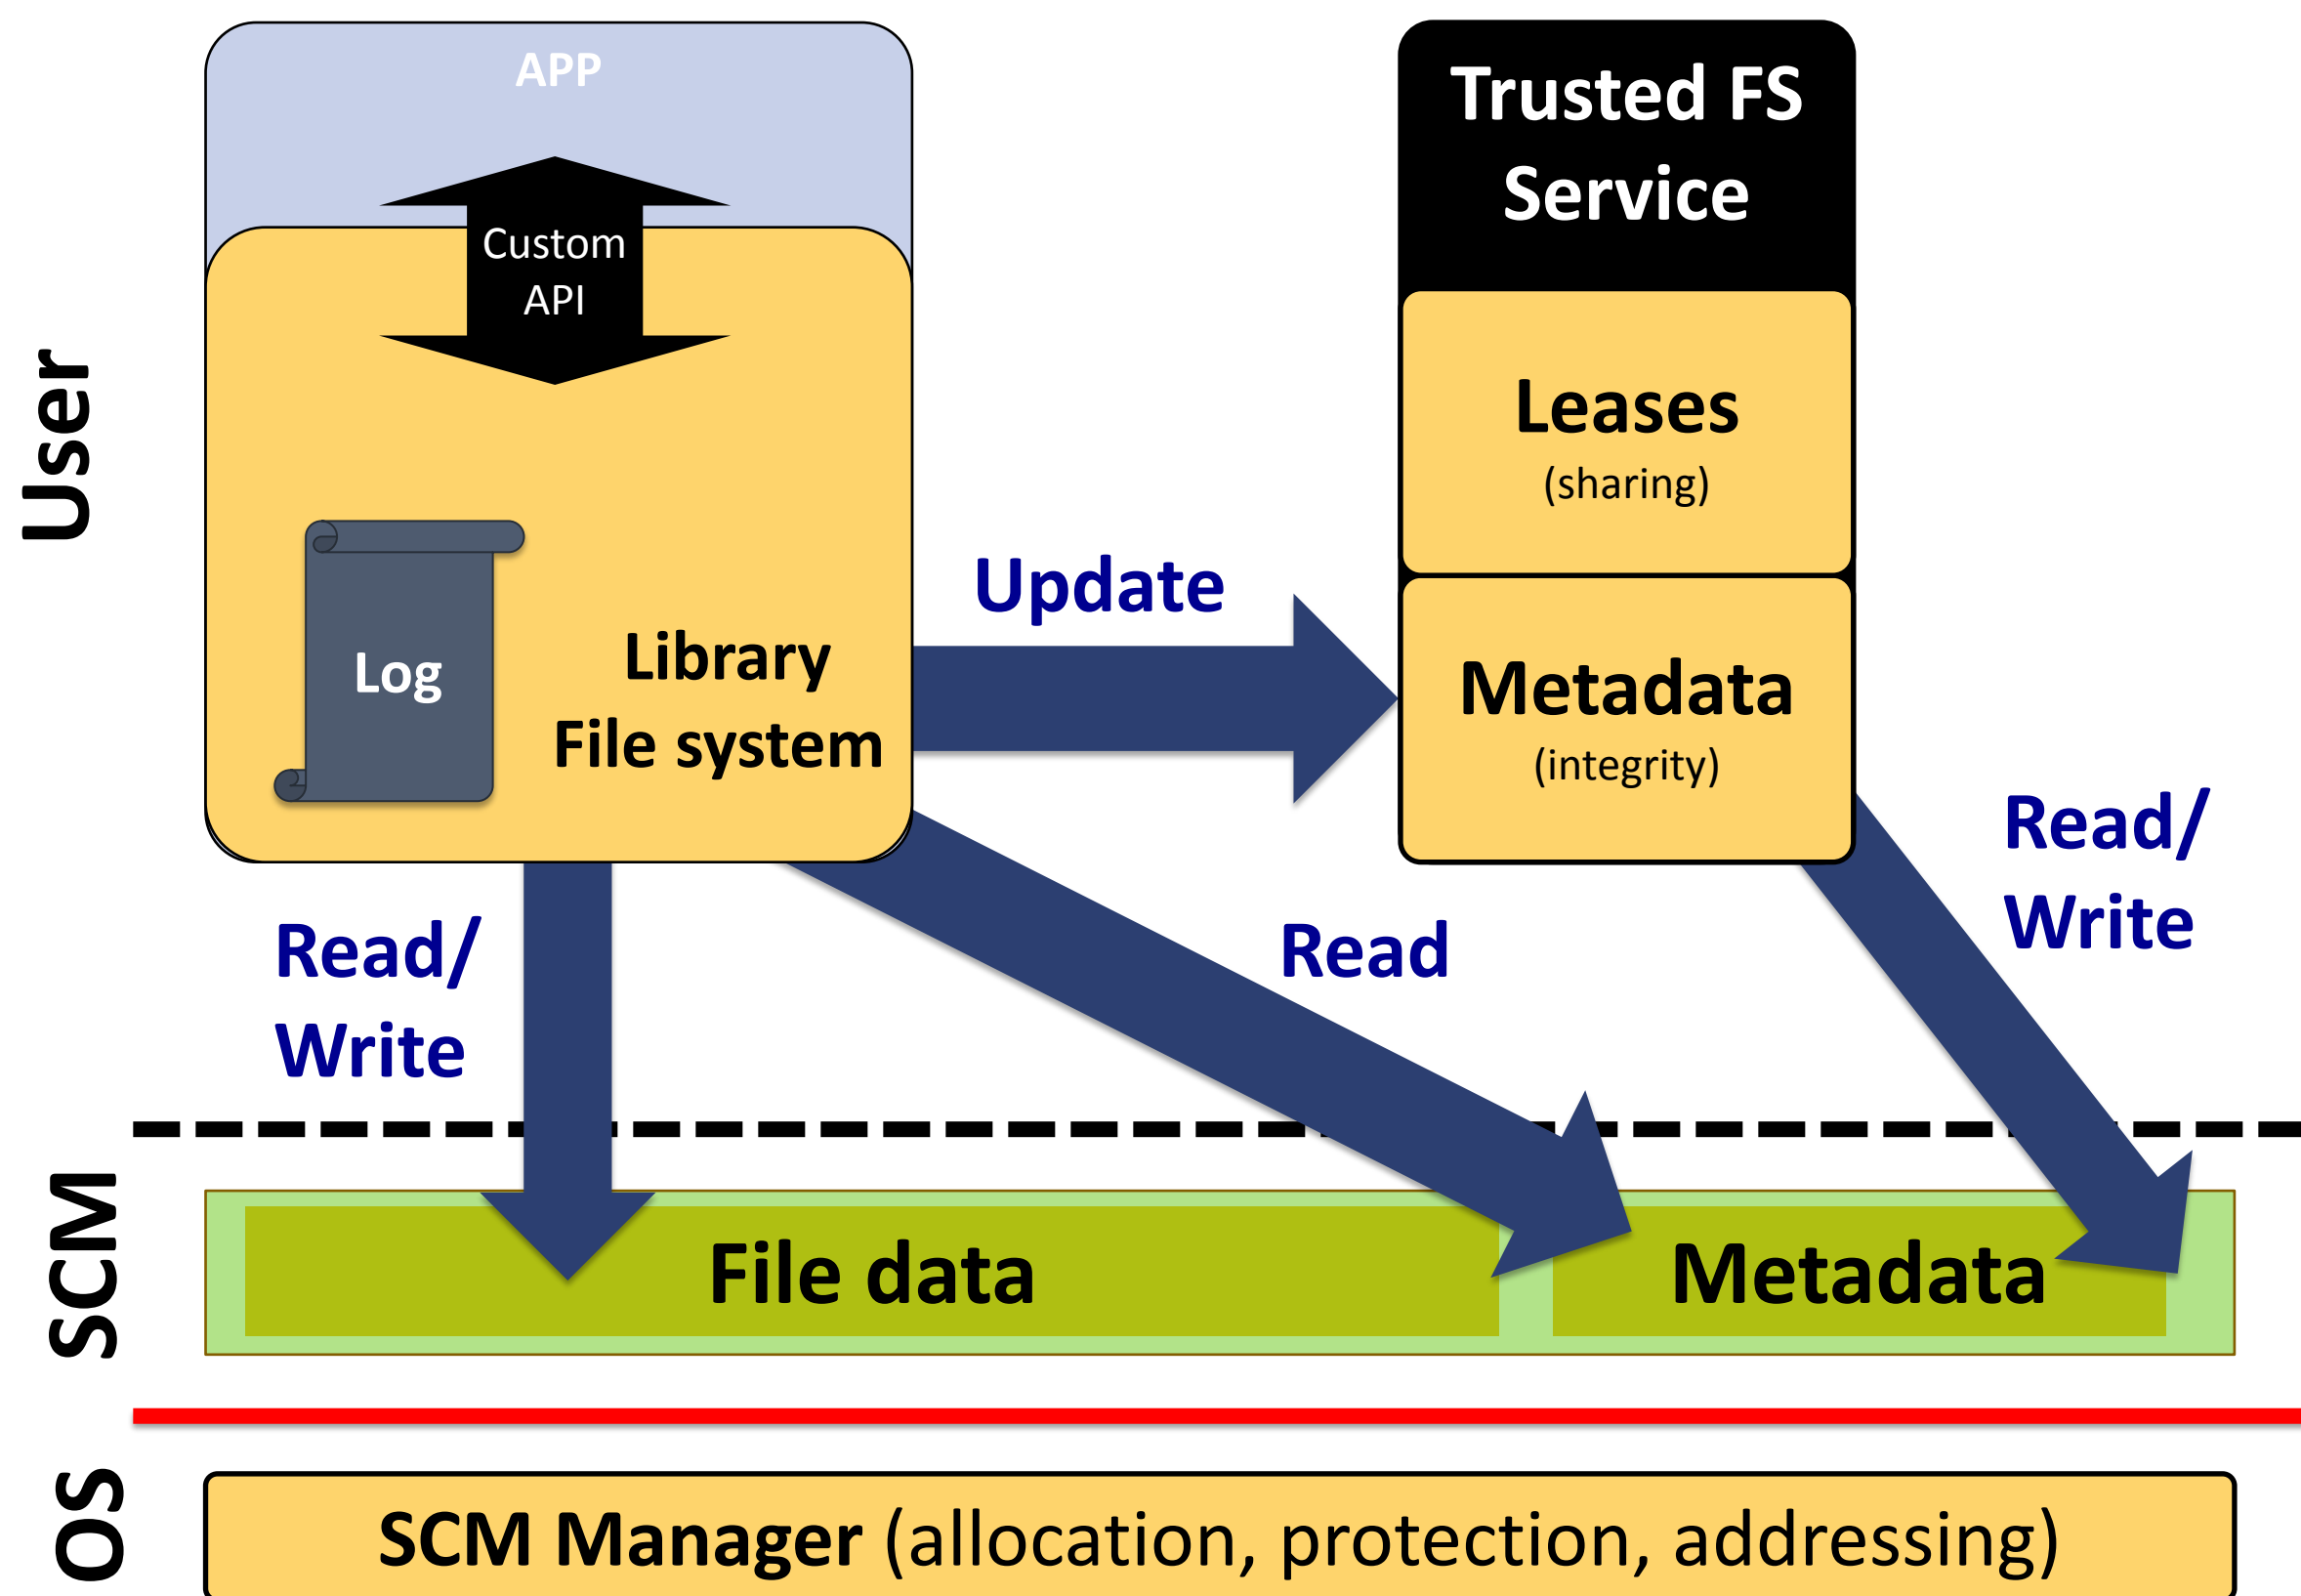
★ **Hierarchical FS impose overhead on fast hardware**
Kernel FS impose overhead and hinder flexible APIs

TOWARDS FLEXIBLE INTERFACES

★ Library Memory File-systems

- Provide **direct access** to SCM
- Enables common operations (**open/read/write/close file data, read meta-data**) without involving trusted centralized entity

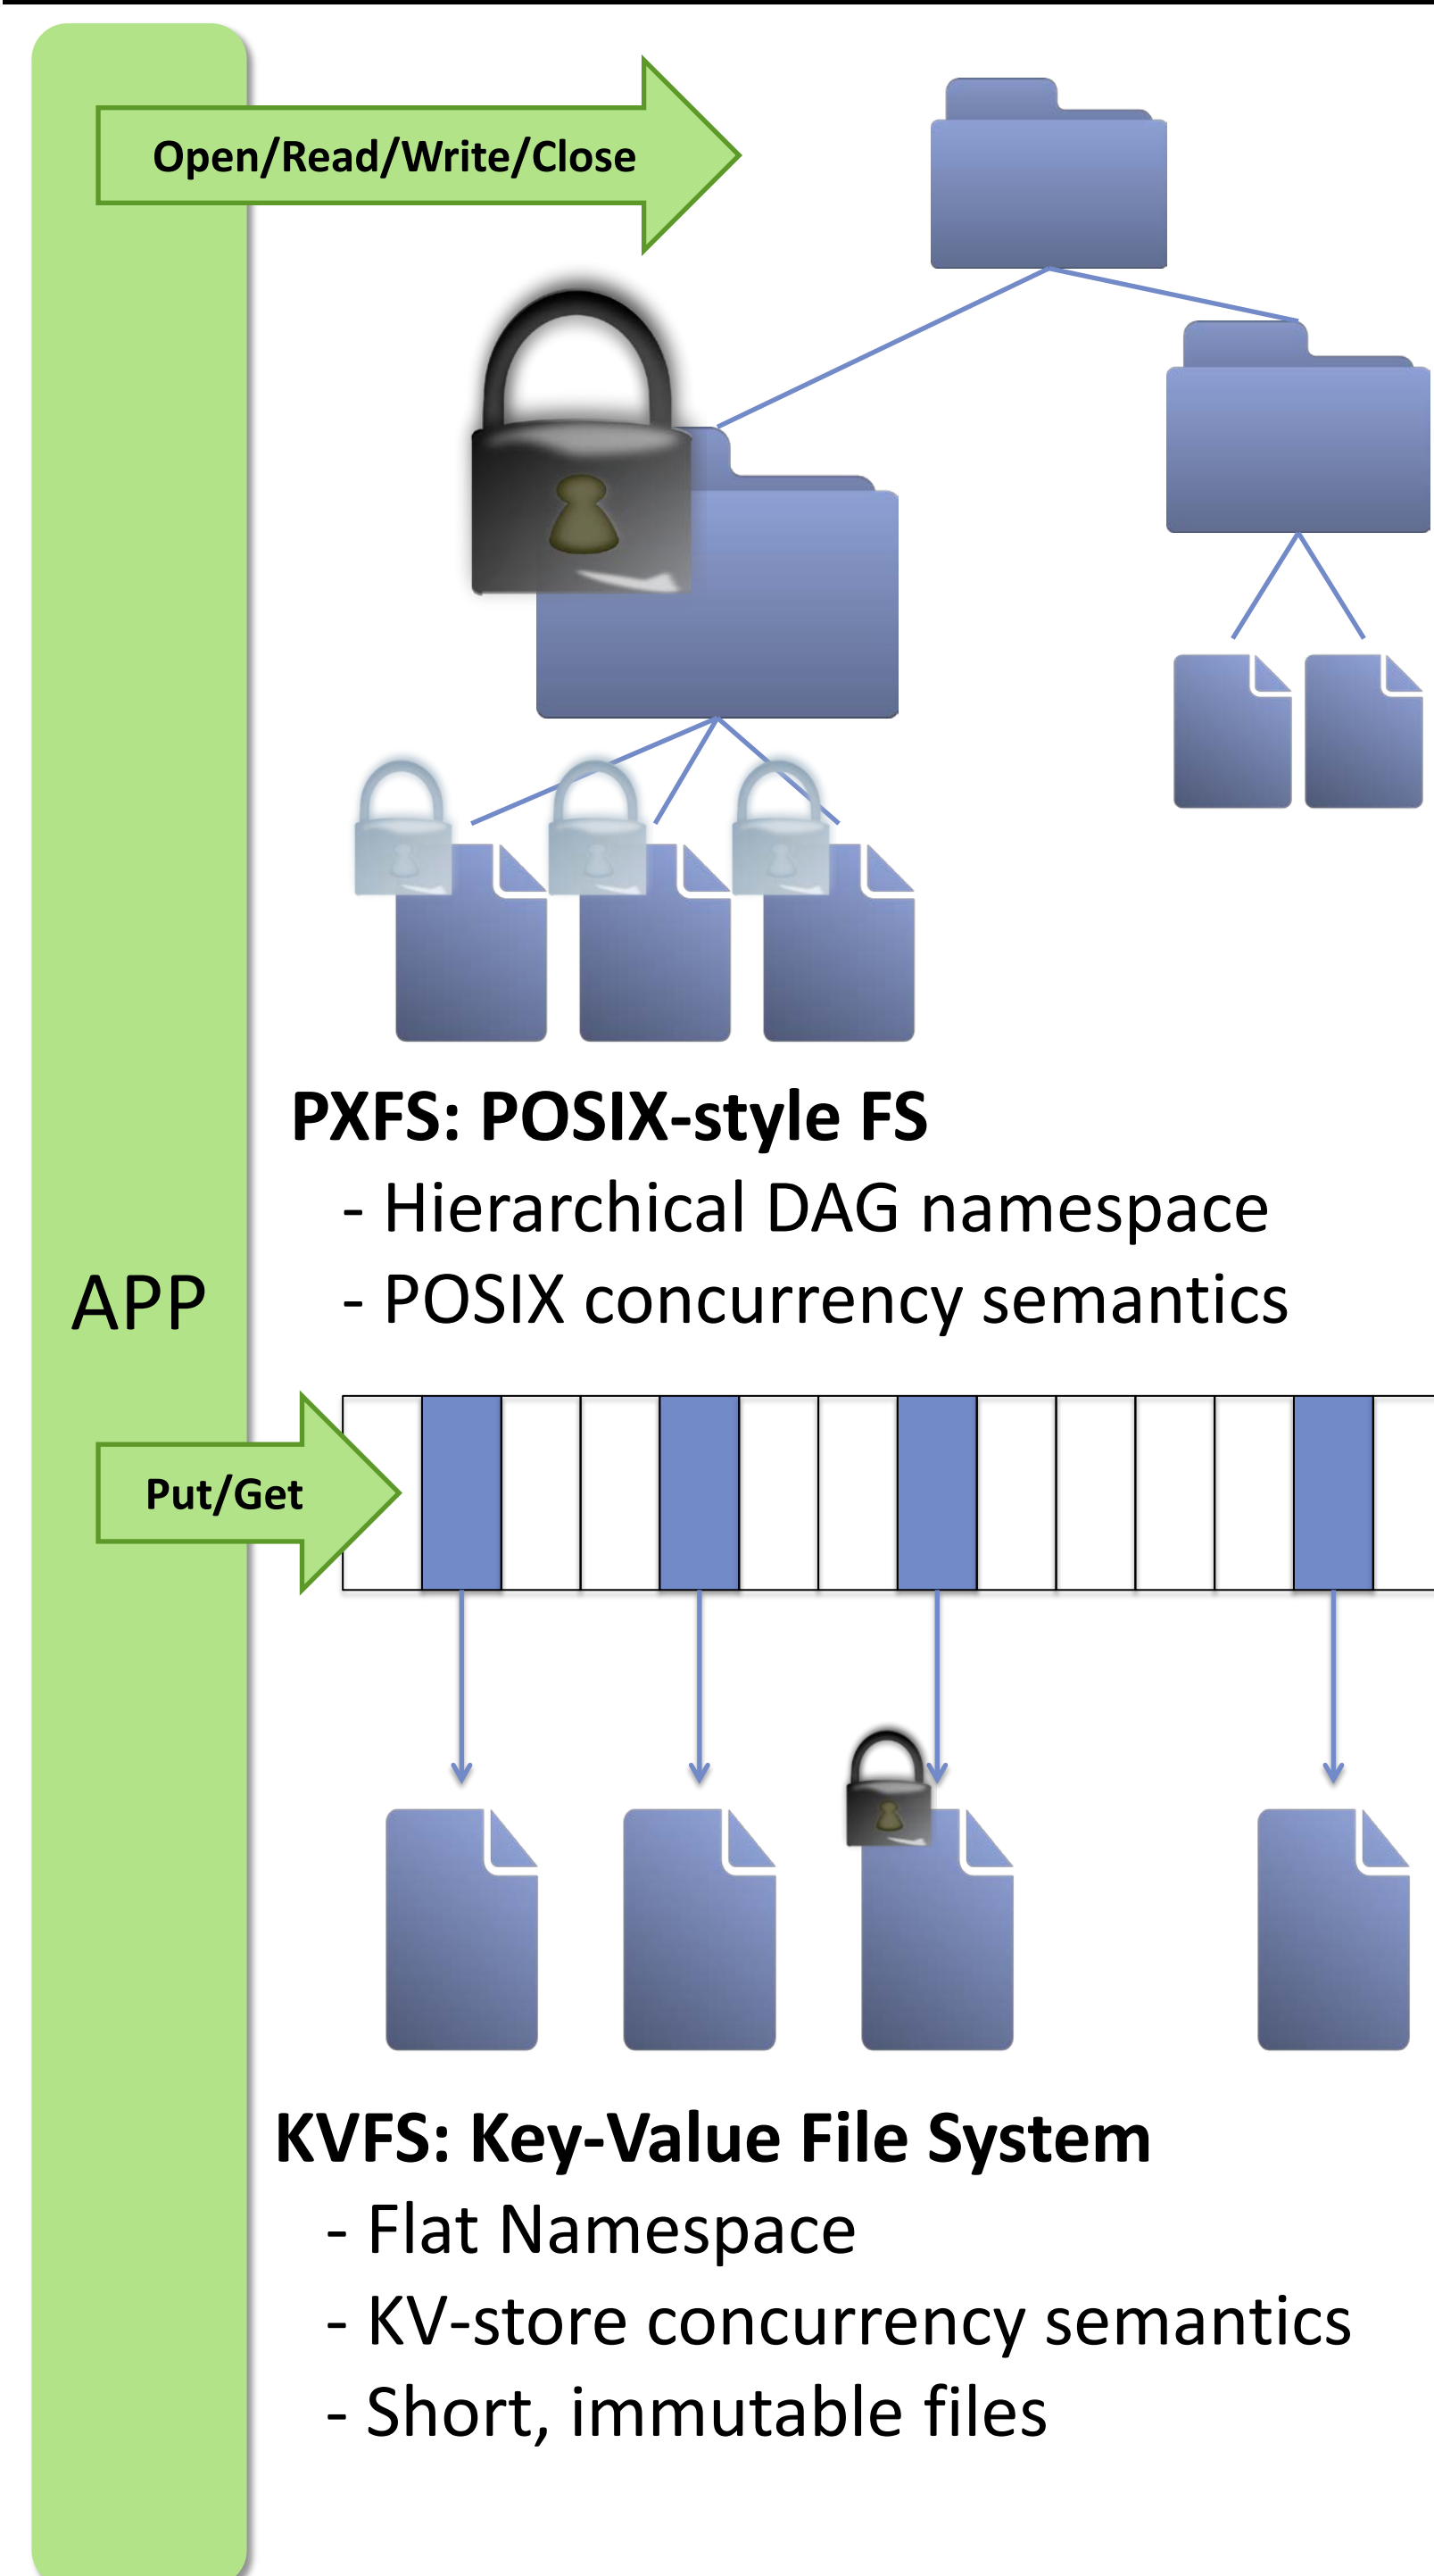
Decentralized Architecture



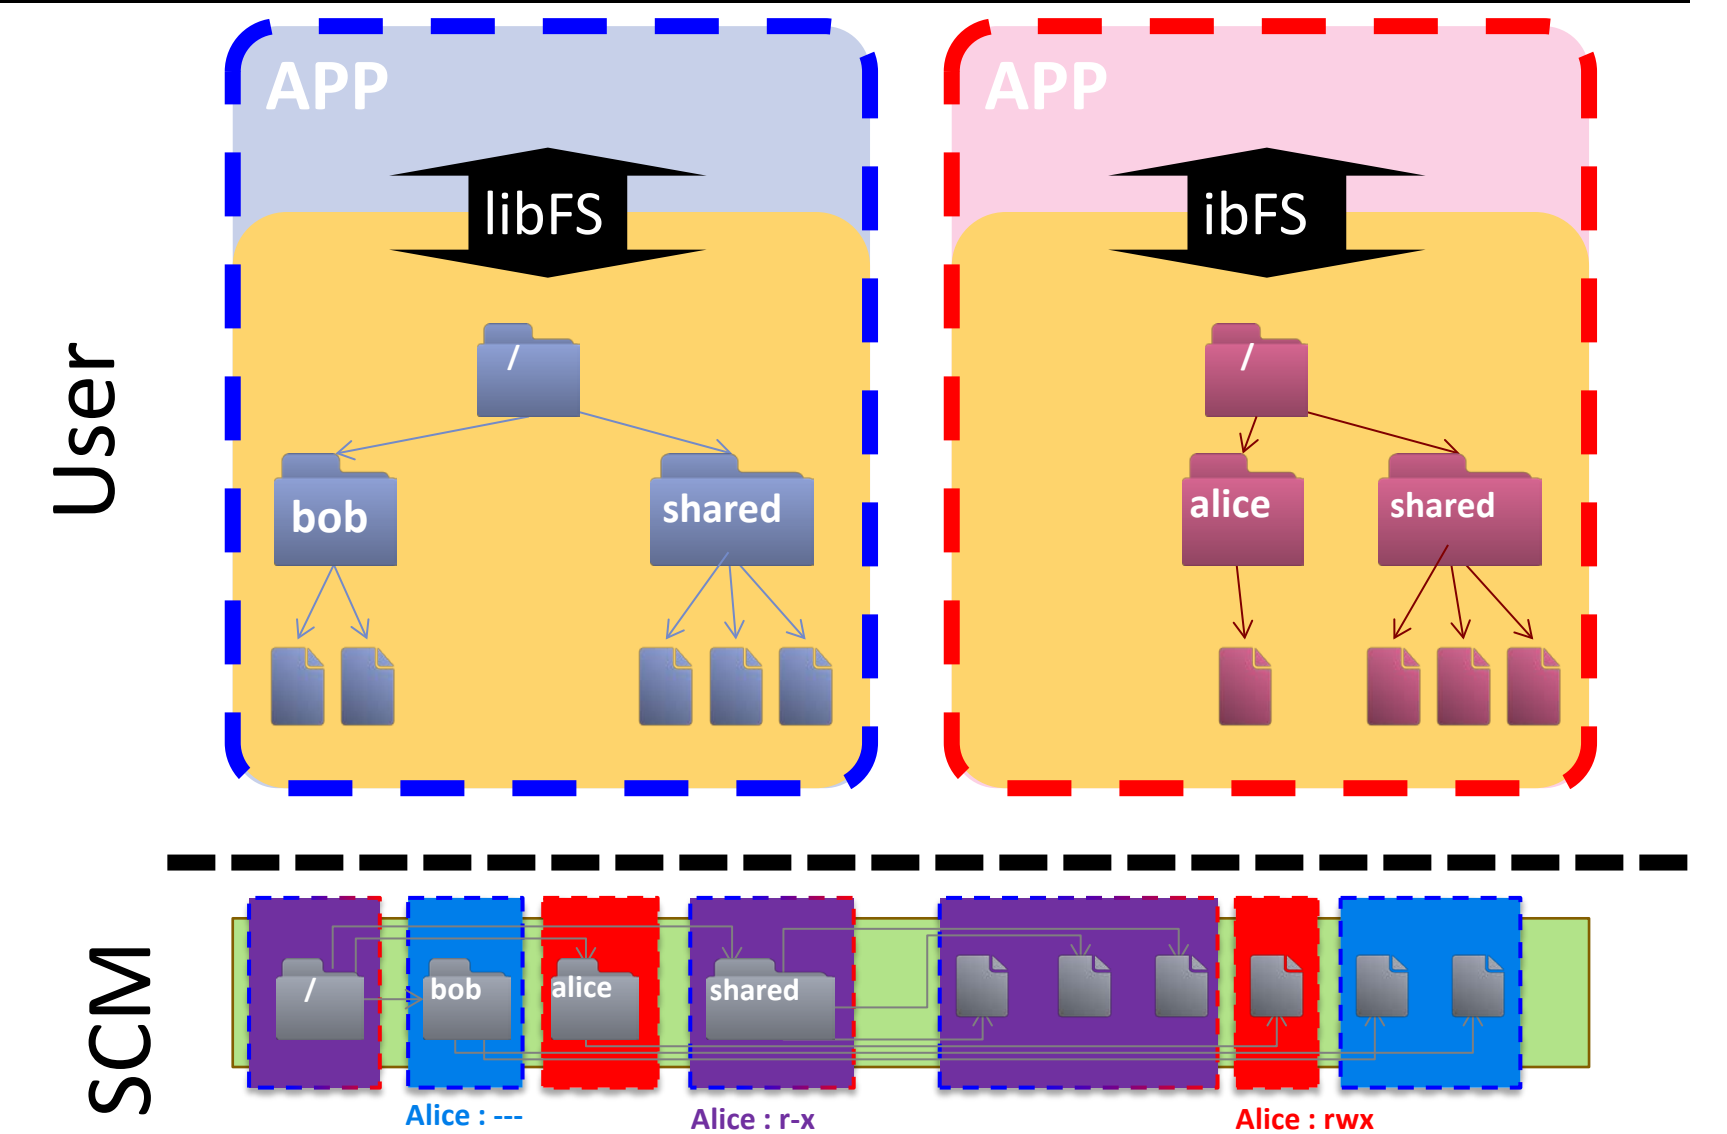
System Components

- **Library File-System (libFS)**
 - Exposes file system API
 - Implements naming and mapping
- **Trusted File-System Service (TFS)**
 - Guarantees integrity of metadata updates
 - Enforces concurrency ctrl
- **SCM Manager**
 - Securely records and enforces resource usage

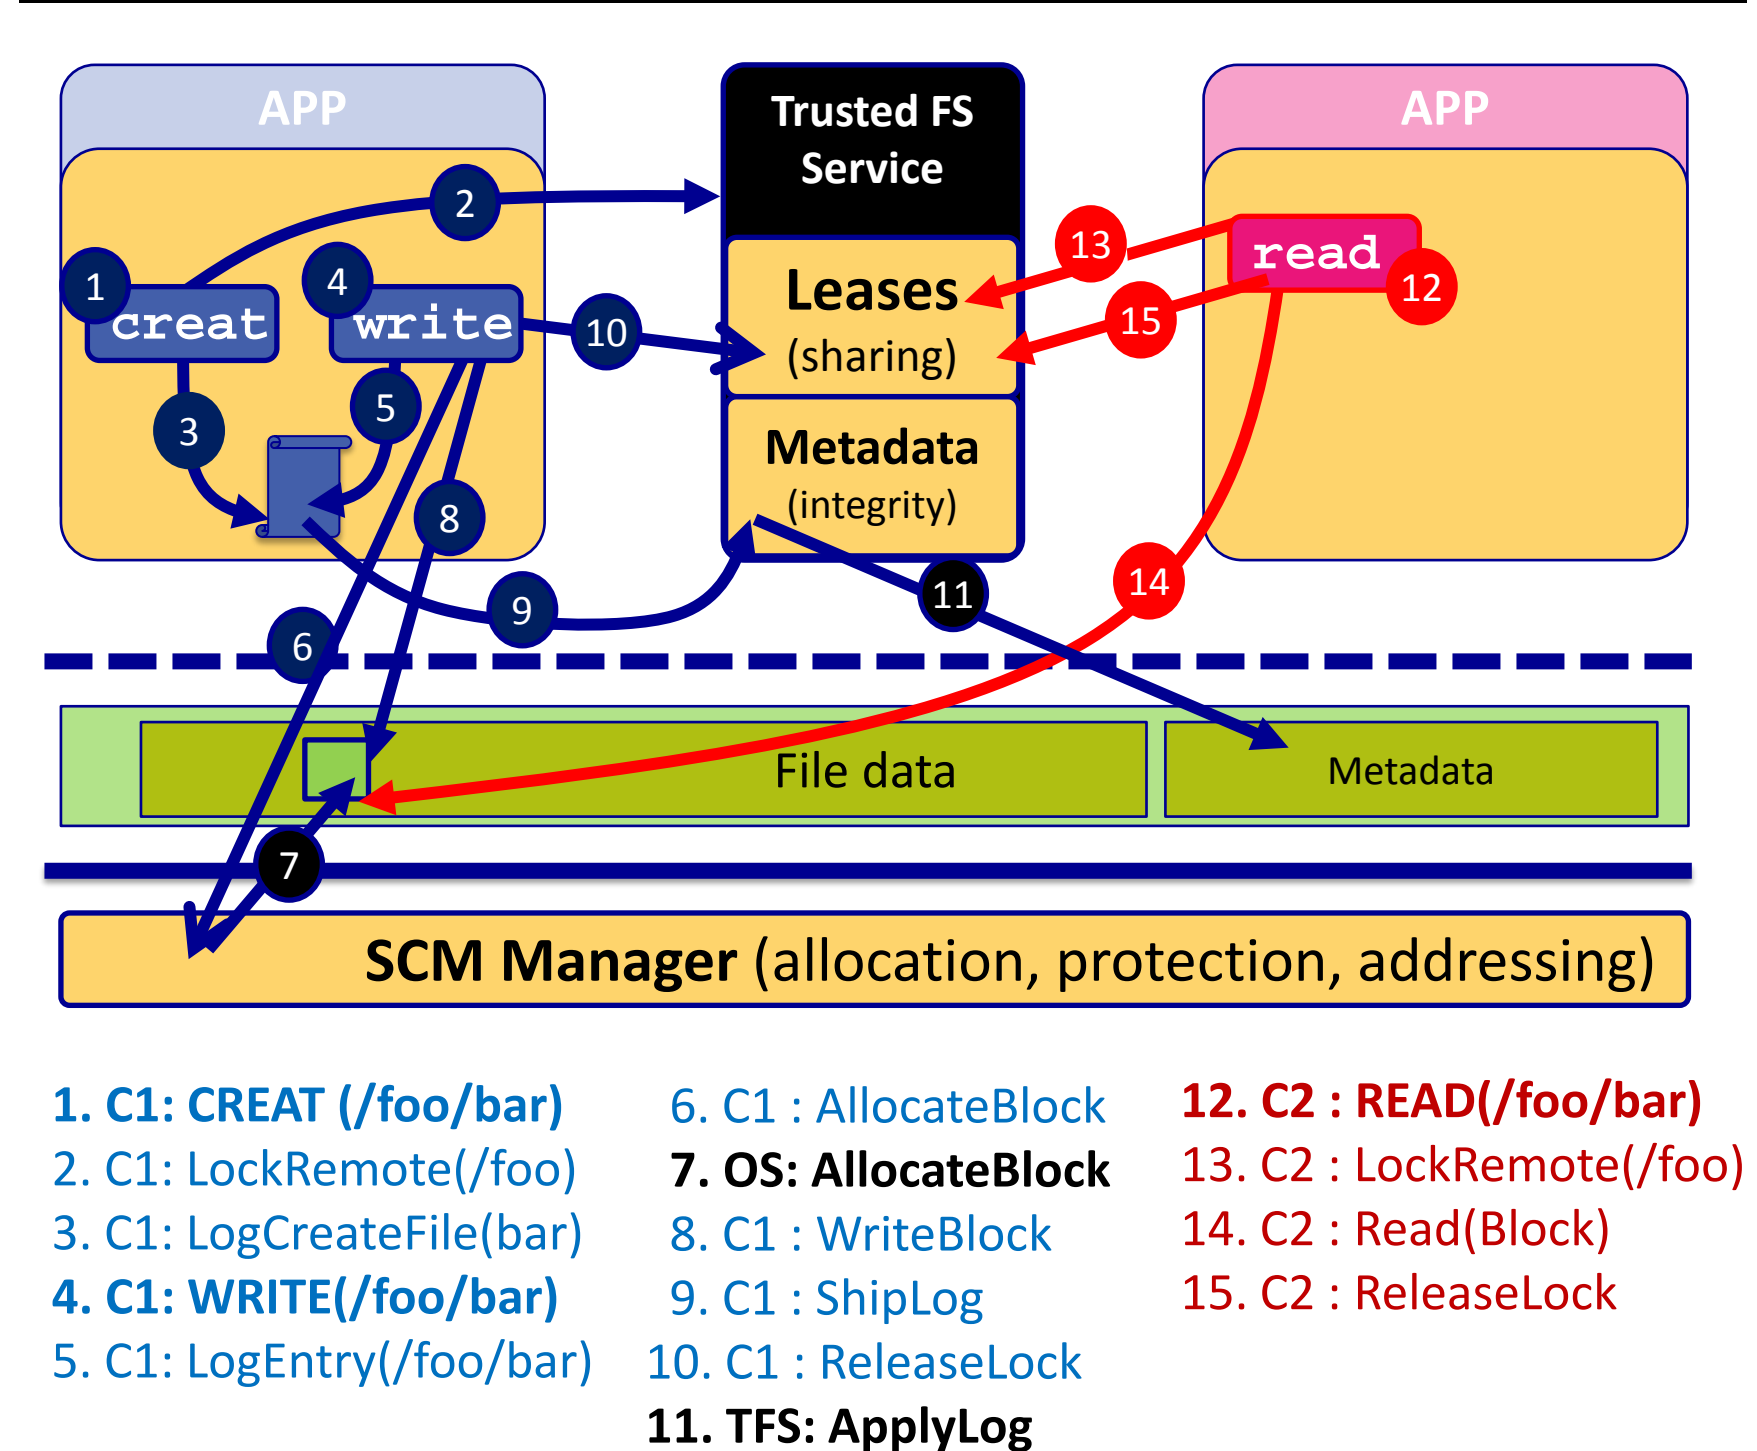
Library File Systems



Enforcing FS Permissions using MMU




Example: The Life of a Shared File




EVALUATION

Client Scalability (DRAM Write + Delay + Read)


File Server



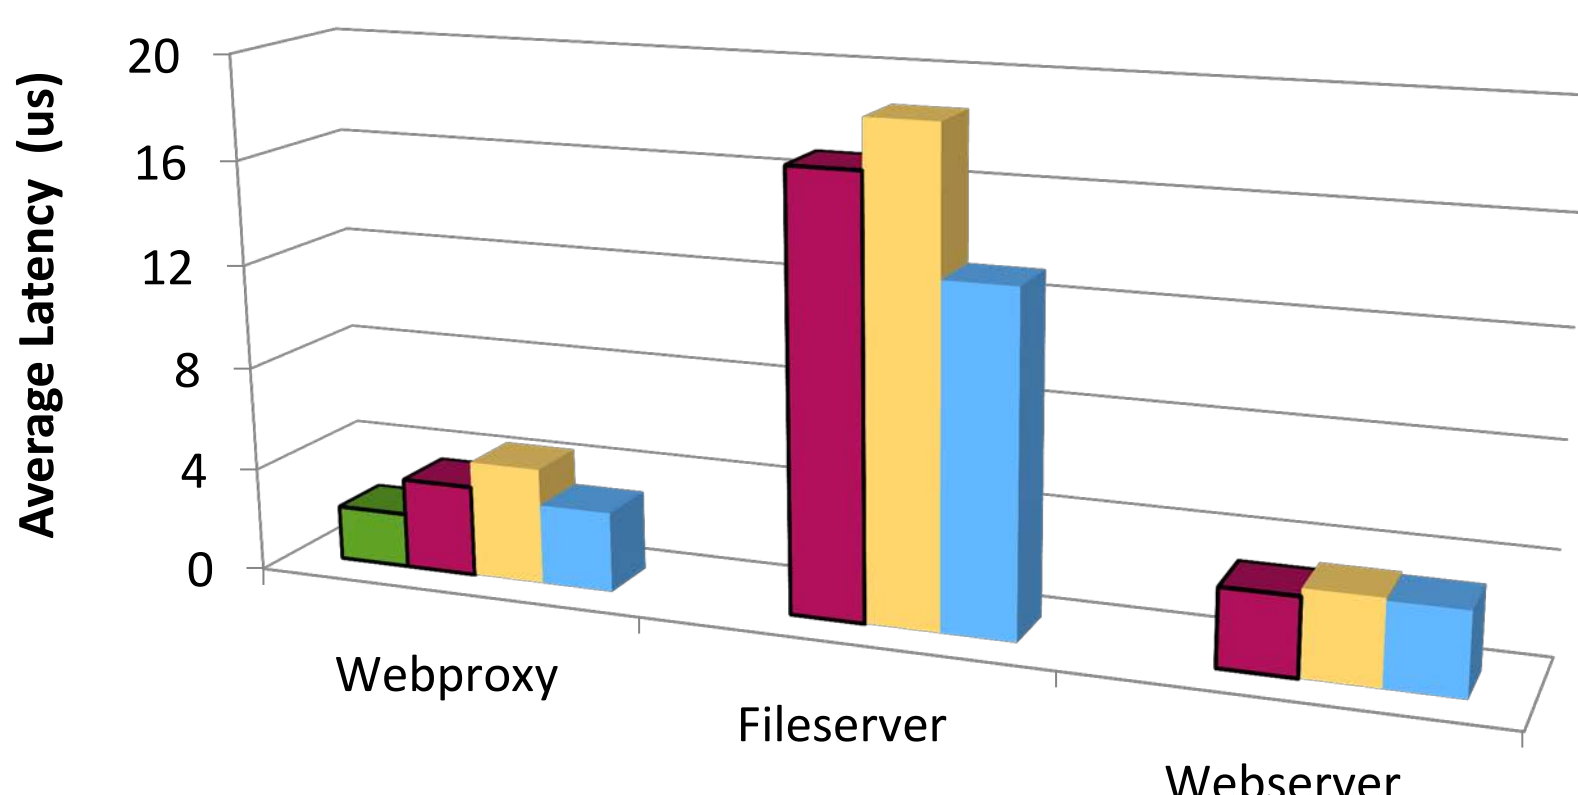
Web Server




Web Proxy



Filebench Latencies



Sensitivity to SCM Performance



★ **Web-proxy workload performs best when there are no overheads of a hierarchical or kernel file system such as ext4 or PXFS !**

Platform:
Intel(R) Xeon(R) E5645 2.40GHz (6 cores with 2 hardware threads per core)
Linux 3.2.2 x86-64
Workload:
Filebench application personalities