

Securing Applications using Operating System Transactions

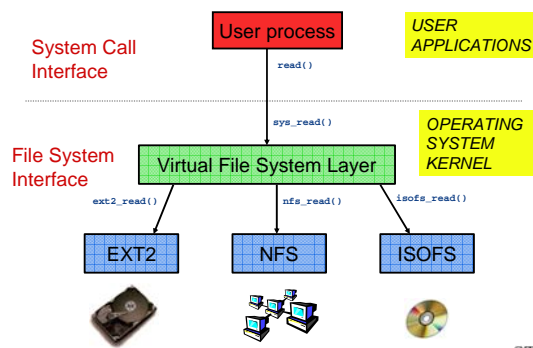
and
File System Security Research Overview

Erez Zadok

File-systems and Storage Laboratory
Stony Brook University
<http://www.fsl.cs.sunysb.edu/>



How File Systems Work

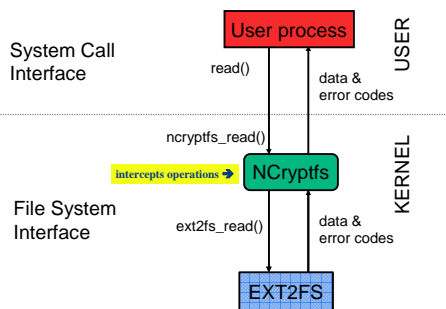


7/14/2006

Zadok - IUCRC CIP organizational workshop

2

File System Interception

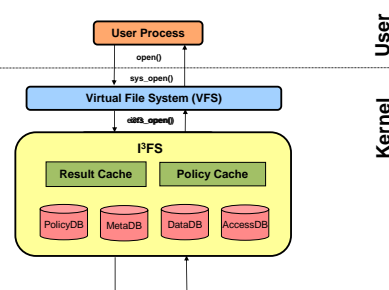


7/14/2006

Zadok - IUCRC CIP organizational workshop

3

Integrity-Checking

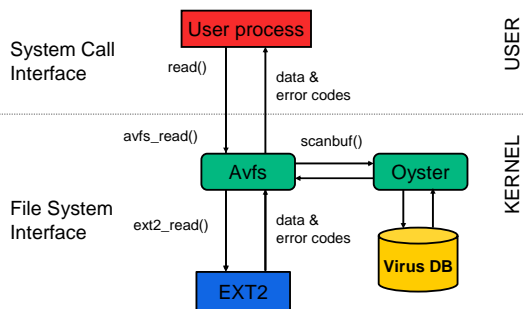


7/14/2006

Zadok - IUCRC CIP organizational workshop

4

Avfs: Anti-Virus File System



7/14/2006

Zadok - IUCRC CIP organizational workshop

5

Secure File System Projects

| File System | Benefit |
|---------------------|--|
| NCryptfs** | Encryption: protect data privacy [Usenix, SISW] |
| i3FS | Detect unauthorized data changes [LISA] |
| Antivirusfs | Transparent anti-virus checking [Usenix Security] |
| Tracefs | Fine-grained monitoring (IDSs) [FAST] |
| Replayfs | Fine-grained replaying (forensics) [FAST] |
| Versions, Snapshots | Checkpoint file state, easy undo and redo (forensics) [FAST] |
| RAIF | Distributed Storage Survivability [ClusterSec**] |
| SDFS | Secure Deletion of files [SISW] |
| E2EC** | End-to-End security using NFSv4 (IBM) |

7/14/2006

Zadok - IUCRC CIP organizational workshop

6

POSIX API Problems

- Guarantees only single system call behavior
 - ◆ but not multiple system calls
 - ◆ OS may not guarantee even that
- System failure leave application state inconsistent on disk
 - ◆ fsck won't help you
- Other system call can interleave
- TOCTTOU security exploits

7/14/2006 Zadok – IUCRC CIP organizational workshop 7 STONY BROOK

Solution? Databases

Applications can use databases:

- ☺ Offer full ACID semantics
- ☹ Cumbersome to use
- ☹ Differing APIs
- ☹ Bloated: SQL, stored procedures, query optimizations, etc.
- ☹ Each application has to be modified

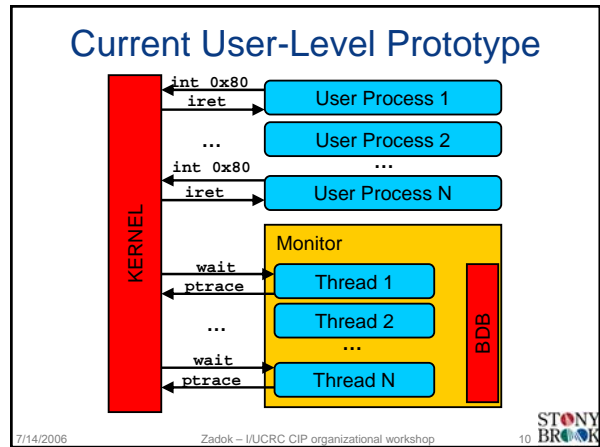
7/14/2006 Zadok – IUCRC CIP organizational workshop 8 STONY BROOK

Our Solution: Berkeley DB

Embed BDB into operating system:

- ☺ Offer full ACID semantics
- ☺ Smaller code base (150k LoC)
- ☺ Maintains POSIX API
- ☺ No need to change applications
- ☺ Easy to use, lightweight
- ☹ Kernel coding is hard

7/14/2006 Zadok – IUCRC CIP organizational workshop 9 STONY BROOK



Example Advantages

- Sample applications we modified
 - ◆ `mail.local` (part of sendmail)
 - reduced code size 450 to 78 (~6x smaller)
 - ◆ `cvs`
 - 347 LoC of "wannabe" code replaced with 3 LoC
 - ◆ `tar`
 - 139 LoC of partially working code replaced with 5 LoC
- Easy to add begin/end transaction calls
 - ◆ Code more reliable
 - ◆ Code more secure
 - ◆ Code easier to audit
- Profiles for legacy application

7/14/2006 Zadok – IUCRC CIP organizational workshop 11 STONY BROOK

