

# Active Data Systems

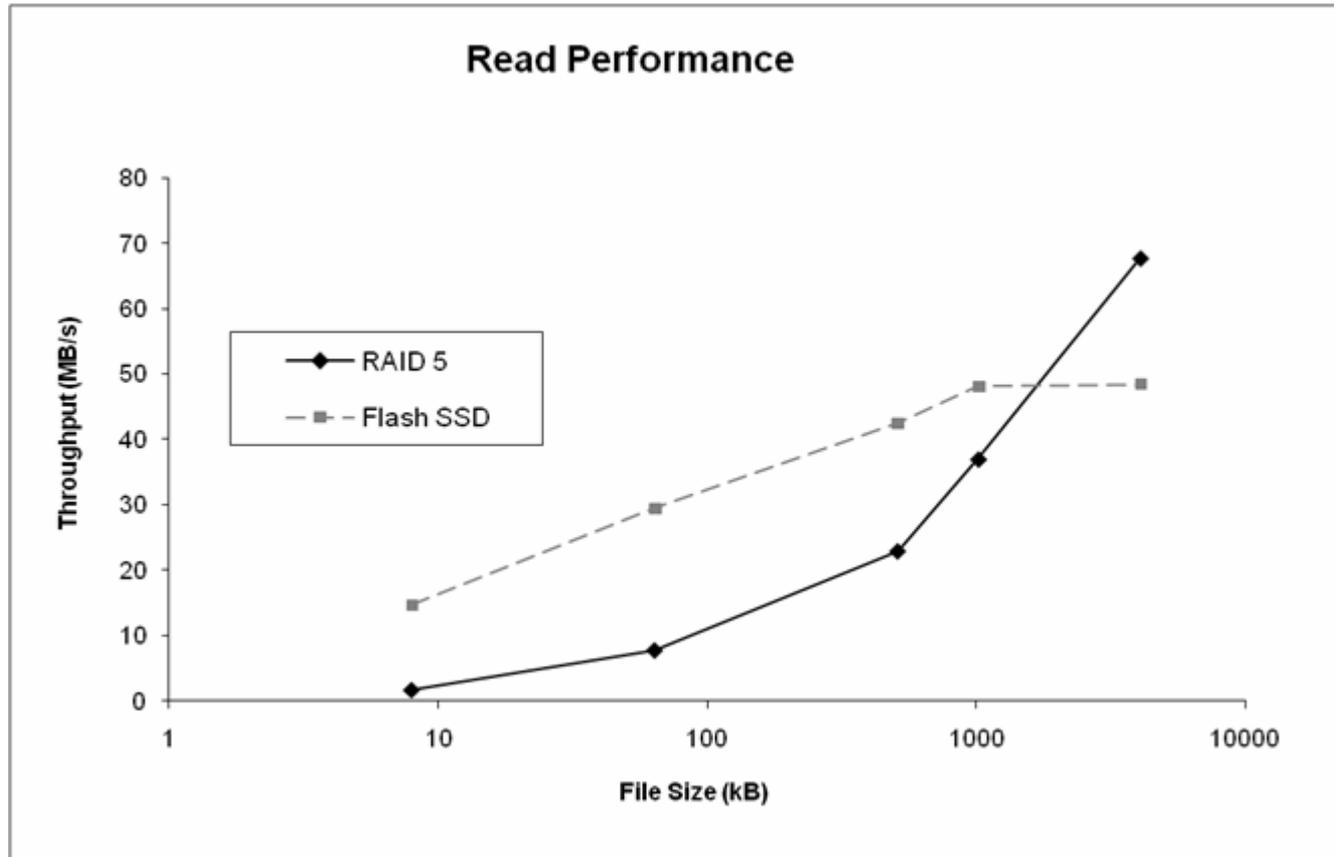
A. L. Narasimha Reddy  
Department of Electrical and Computer Engineering  
Texas A & M University

Students: Sukwoo Kang (now at IBM Almaden)  
John Garrison, Xiaojian Wu

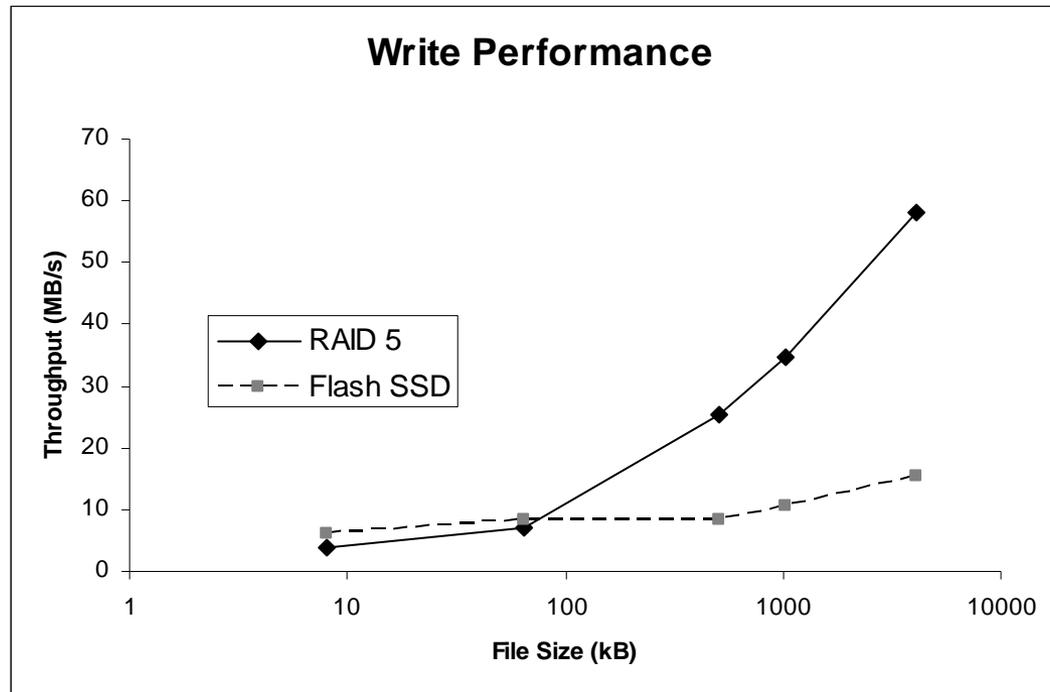
# Storage Management Across Diverse Devices

- **Flash storage becoming widely available**
  - More expensive than hard drives
  - Faster random accesses
  - Low Power consumption
- **In Laptops now**
- **In hybrid storage systems soon**
- **Manage data across Different Devices**
  - Match application needs to device characteristics
  - Optimize for performance, power consumption

# Flash vs. RAID5 Read Performance



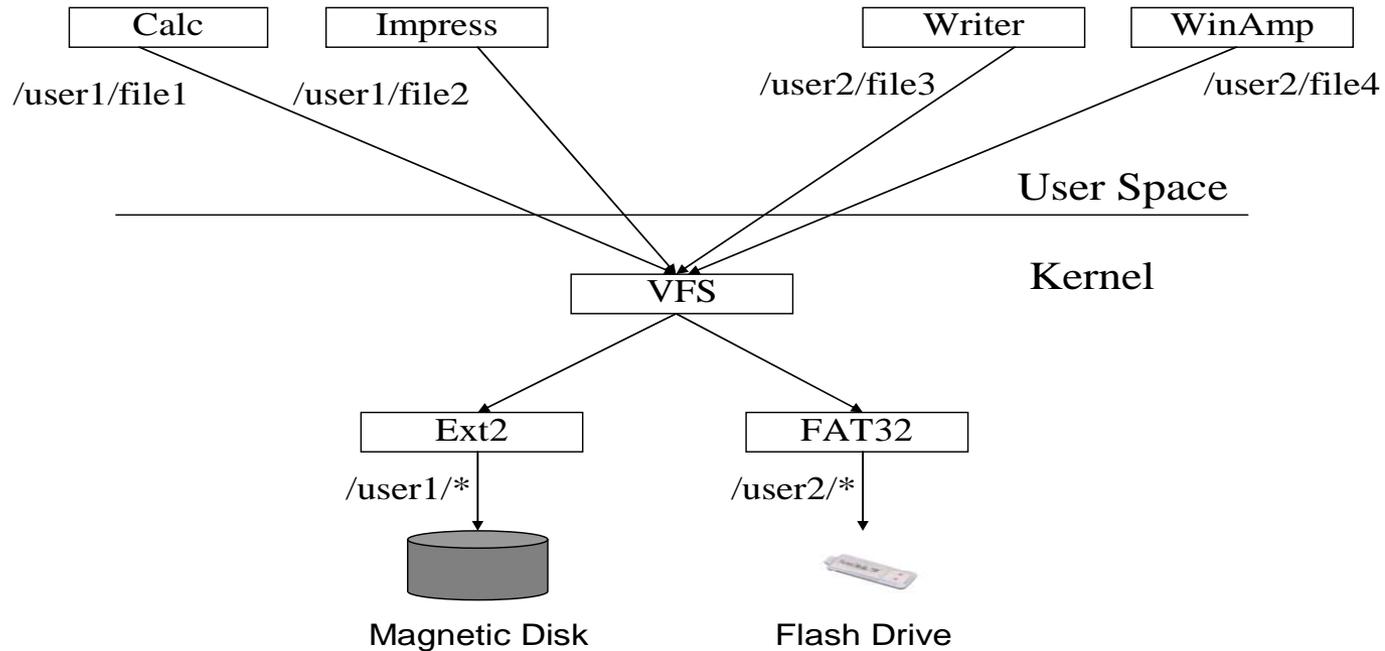
# Flash vs. RAID5 Write Performance



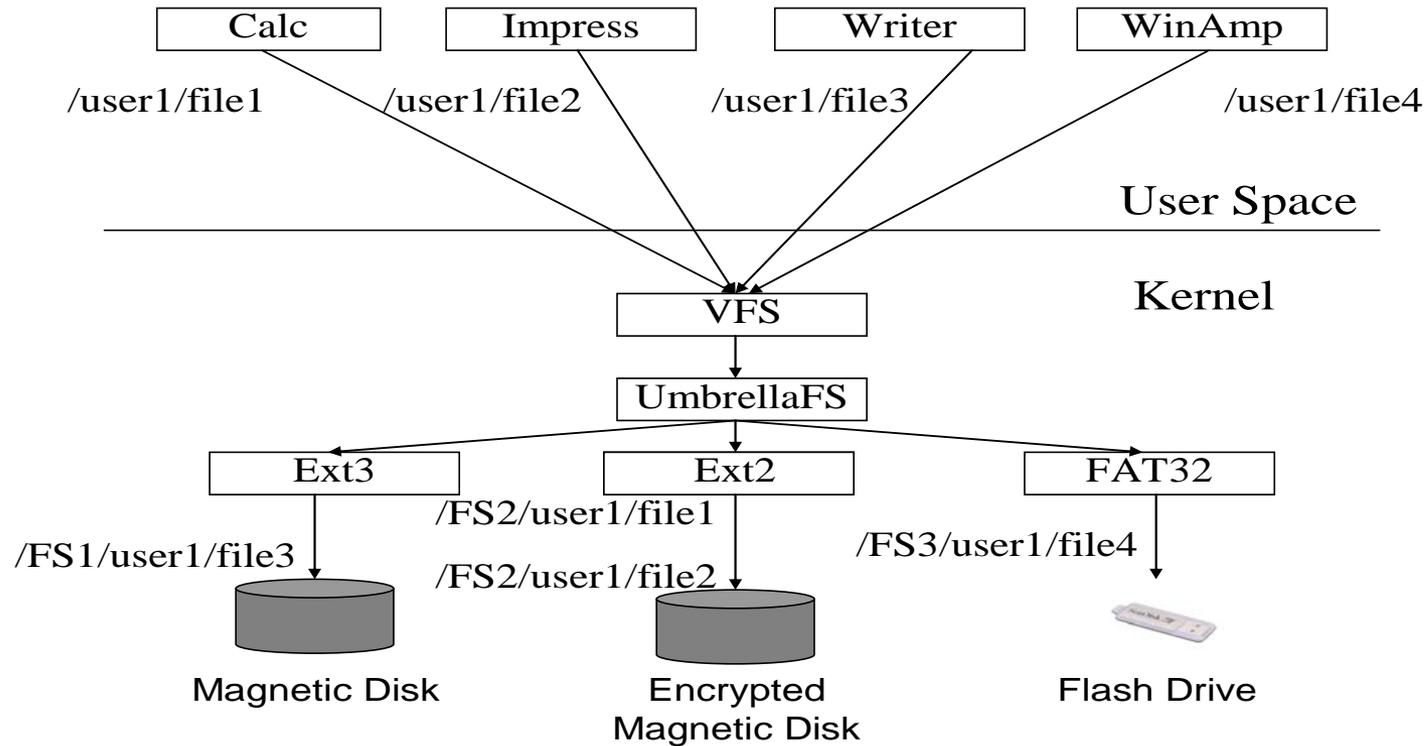
# Motivation

- **Can we provide different storage options for different files for a single user?**
  - /user1/file1 → storage system 1, /user1/file2 → storage system 2...

# Normal File System Architecture

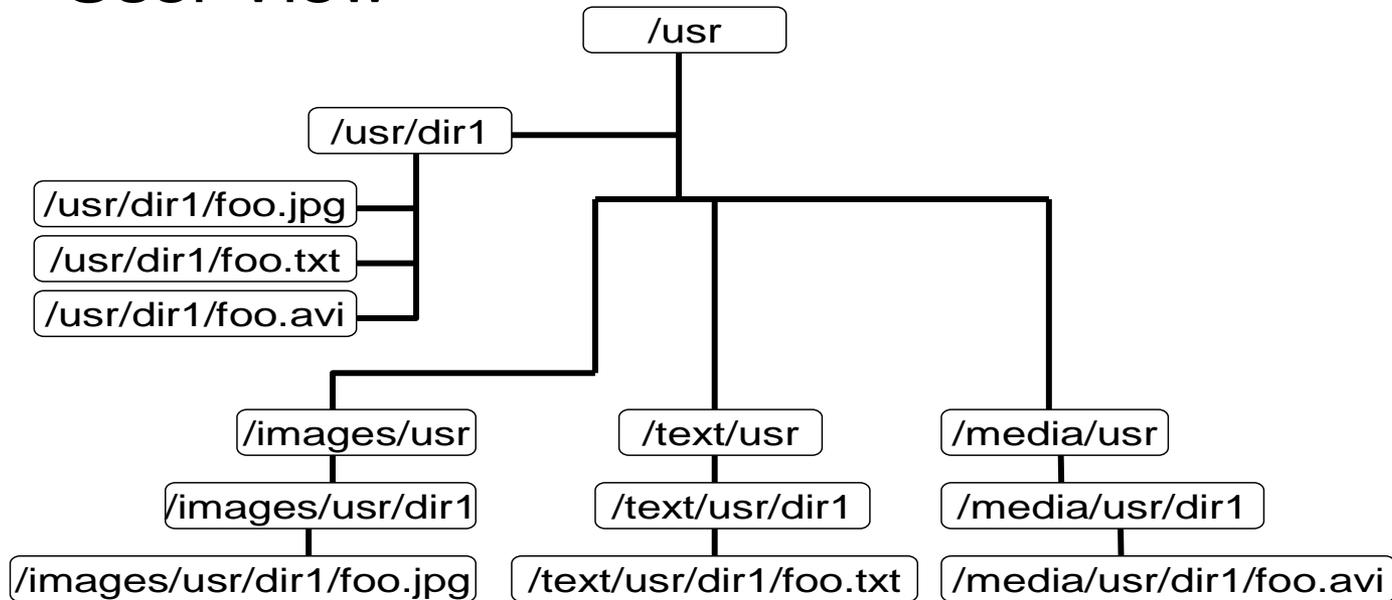


# Umbrella File System



# Example Data Organization

## User View



## Underlying data organization

# Policy Based Storage

- **User or System administrator Choice**
  - Allow different types of files on different devices
  - Reliability, performance, power consumption
- **Layered Architecture**
  - Leverage benefits of underlying file systems
  - Map applications to file systems and underlying storage
- **Policy decisions can depend on namespace and metadata**
  - Example: Files not touched in a week → slow storage system

# Rules Structure

- **Provided at mount time**
- **User specified**
- **Based on inode values (metadata) and filenames (namespace)**
- **Provides array of options**

# Umbrella File System

- **Sits under VFS to enforce policy**
- **Policy enforced at open and close times**
- **Policy also enforced periodically (less often)**
- **UmbrellaFS acts as a “router” for files**
  - Not only based on namespace, but also metadata

# Inode Rules Structure

Rule	Inode/ Filename	Field	Match	Value	Branch
1	Inode	file permissions	=	Read Only	/fs1, /fs2
2	Filename	n/a	n/a	n/a	n/a
3	Inode	file creation time	>=	8:00 am, August 3 <sup>rd</sup> , 2007	/fs2
4	Inode	file length	<	20 KB	/fs3
...					

# Inode Rules

- **Provide in order of precedence**
- **First match**
- **Compare inode value to rule**
  - At file creation some inode values indeterminate
  - Pass over those rules

# Filename Rules Structure

<b>Rule</b>	<b>Match String</b>	<b>Branch</b>
<b>1</b>	<b>/*.avi</b>	<b>/fs2,/fs1</b>
<b>2</b>	<b>/home/*.txt</b>	<b>/fs1</b>
<b>3</b>	<b>/home/jgarrison/*</b>	<b>/fs3</b>
<b>...</b>		

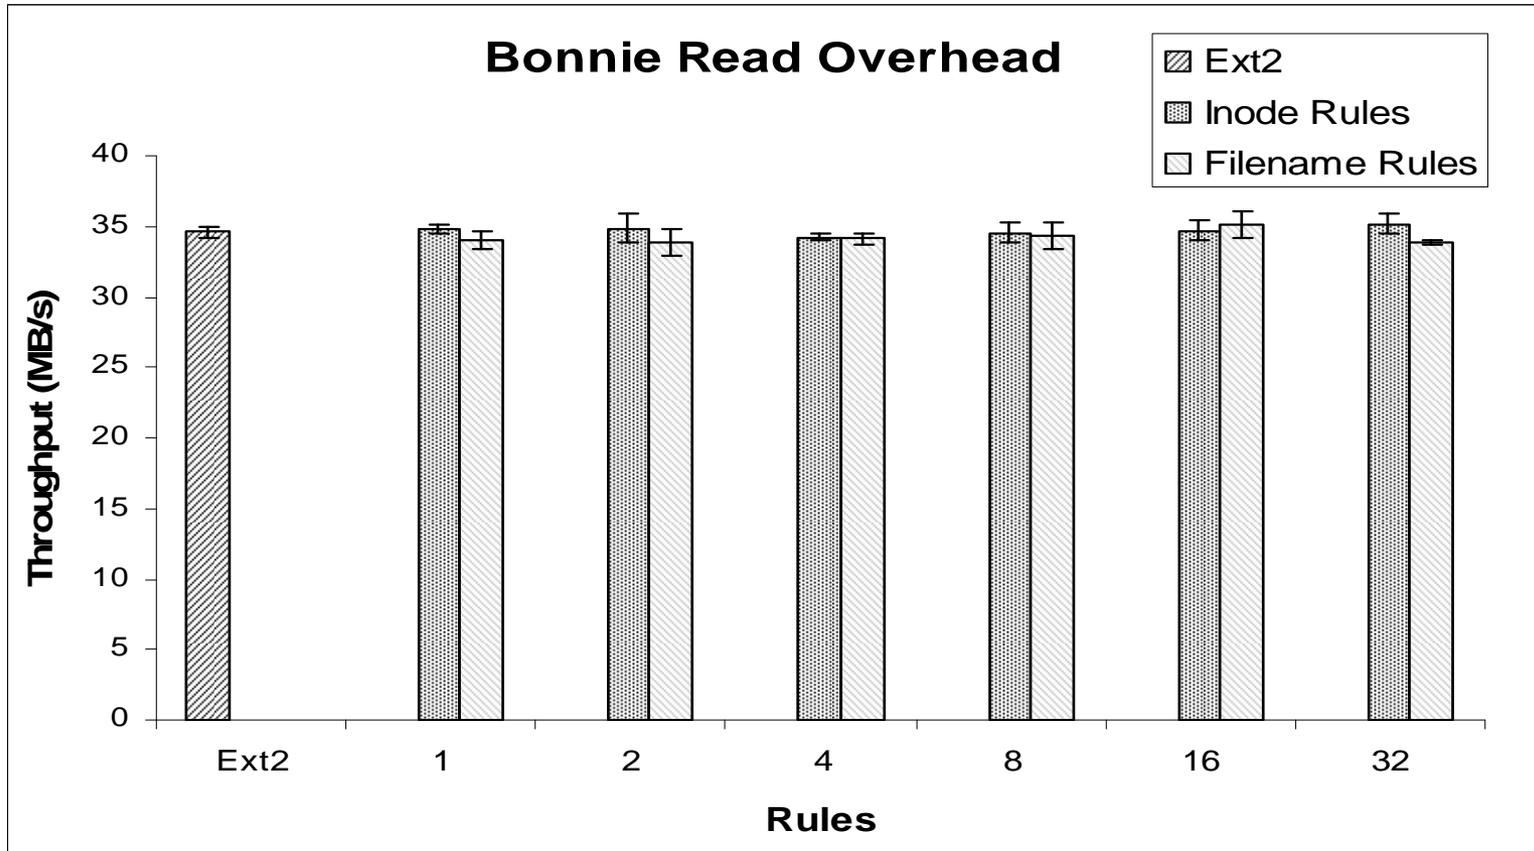
# Filename Rules

- **Once first filename rule triggered, all checked**
- **Similar to longest prefix matching**
- **Double index based on**
  - Path matching
  - Filename matching
- **Example:**
  - Rules: /home/\*/\*.bar, /home/jgarrison/foo.bar
  - File: /home/jgarrison/foo.bar
  - File matches second rule more closely (3 path length and 7 characters of file name vs. 3 path length and 4 characters of file name)

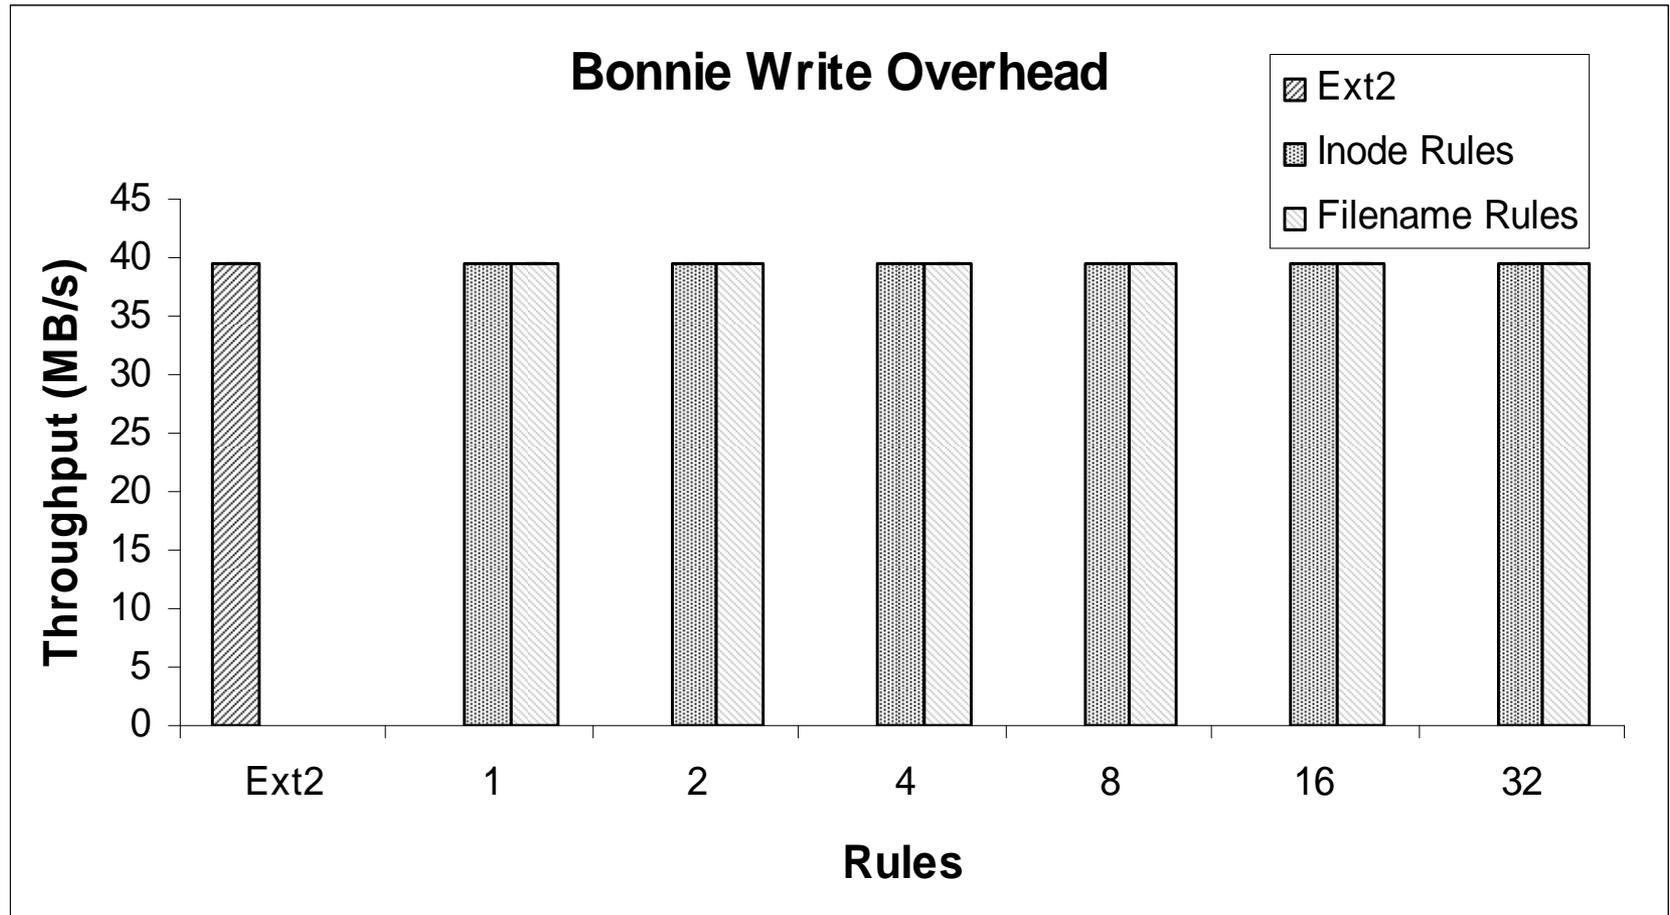
# Evaluation

- **Overhead**
  - Throughput
  - CPU Limited
  - I/O Limited
- **Example Improvement**

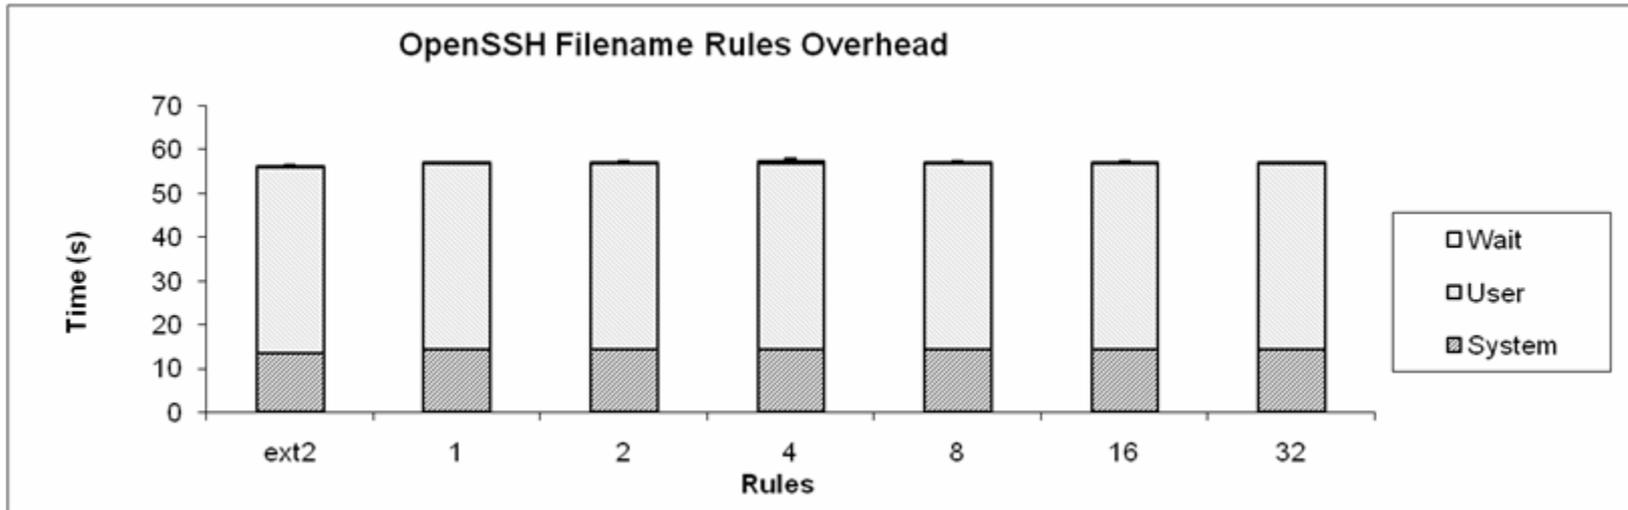
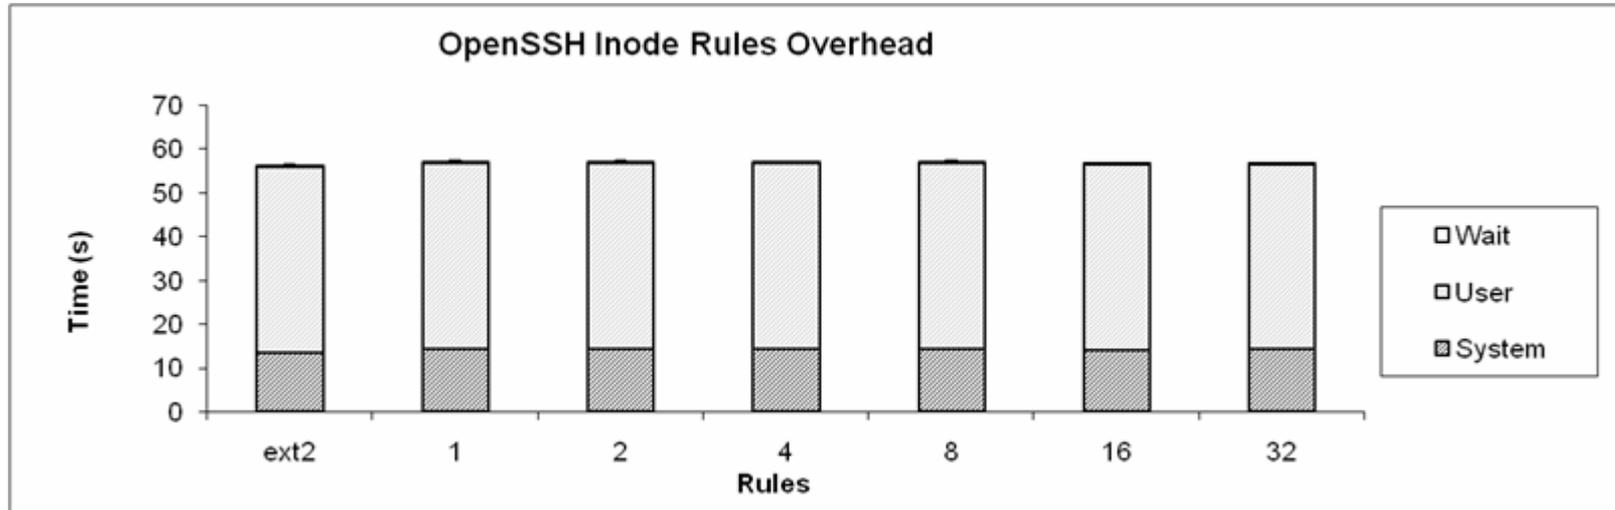
# UmbrellaFS Overhead



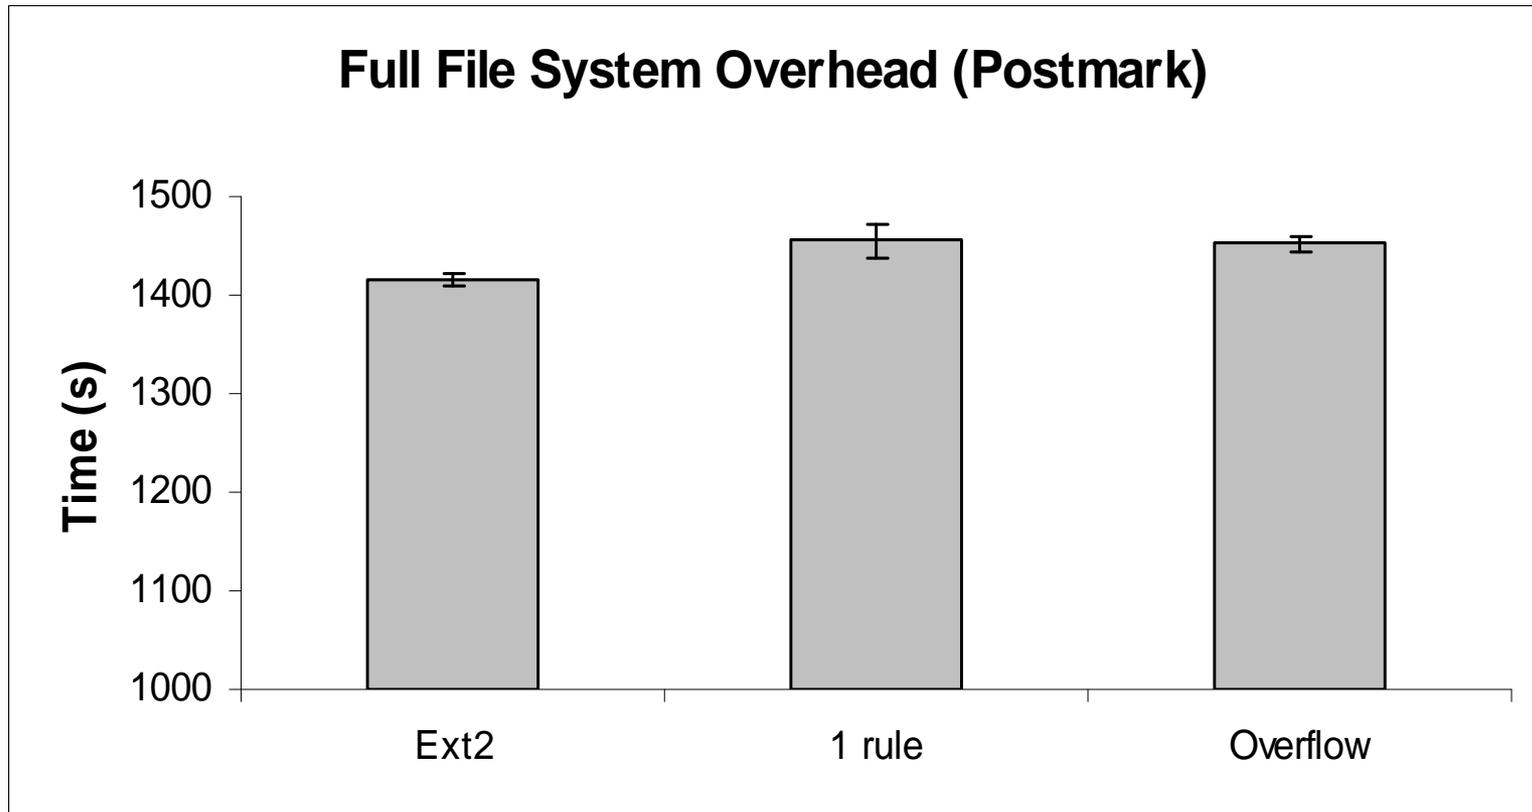
# UmbrellaFS overhead



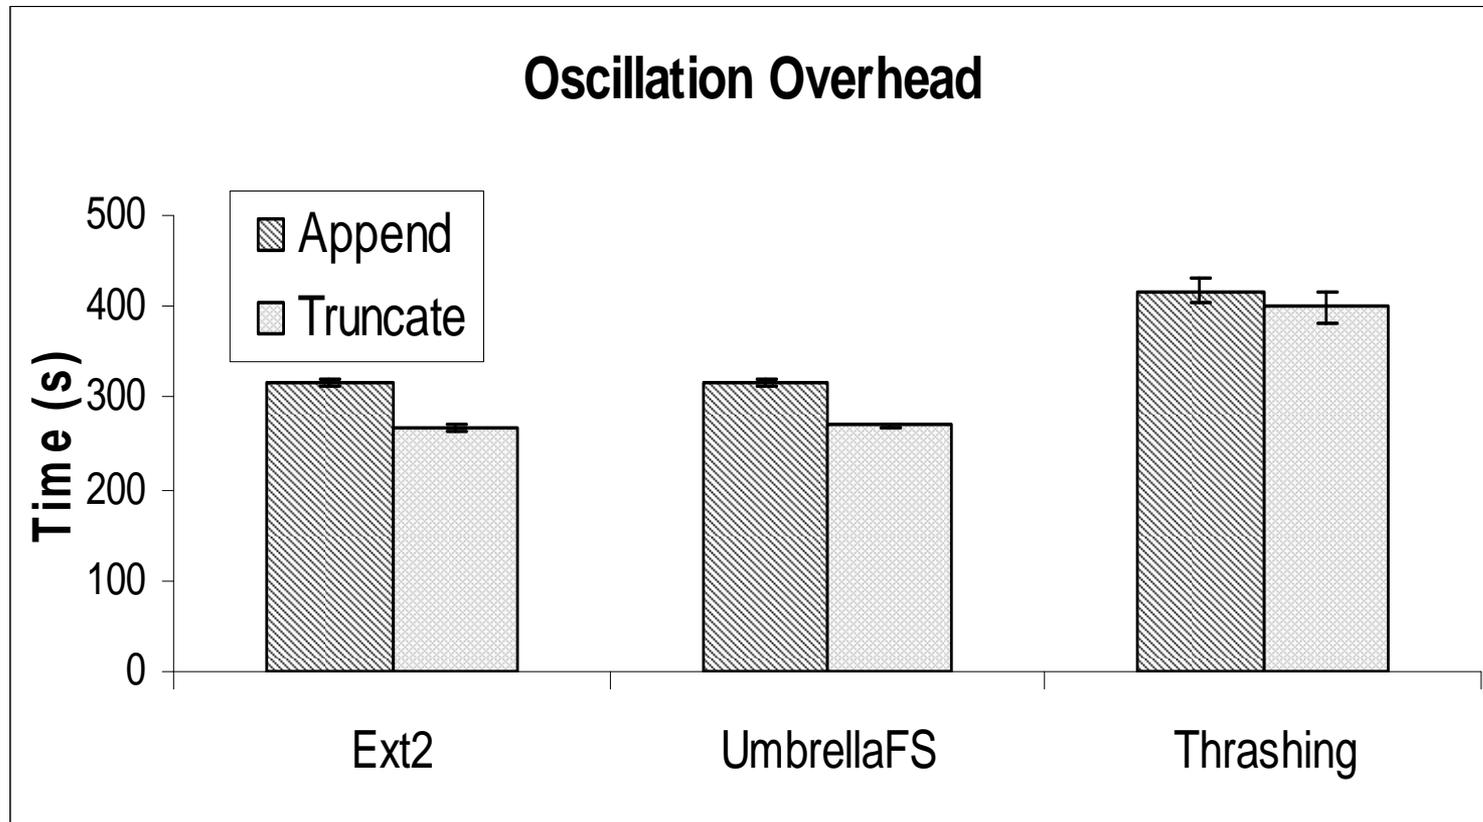
# CPU Limited Benchmarks



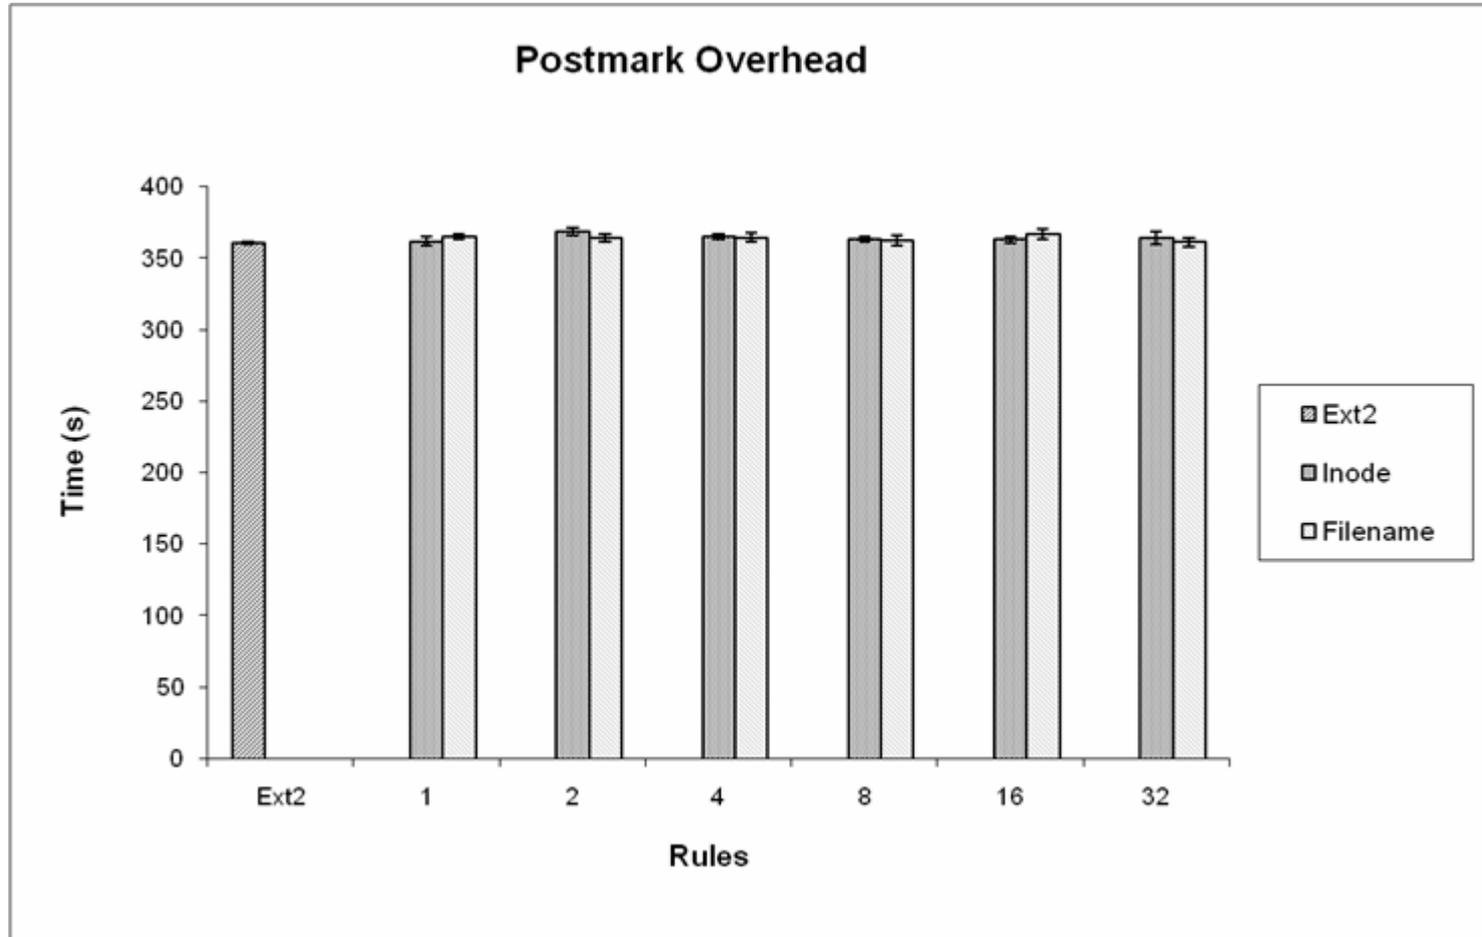
# Overflow (2 file systems) overhead



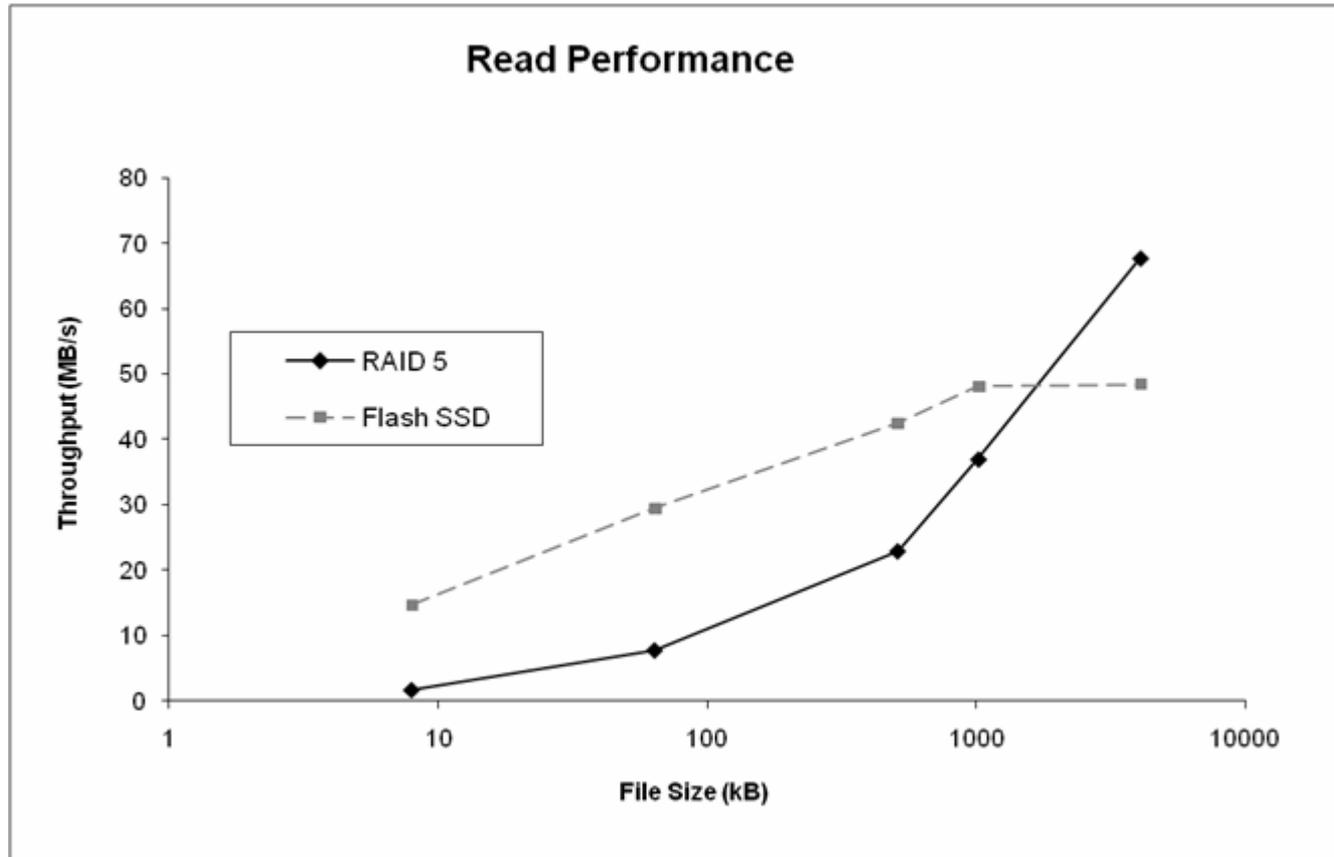
# Oscillation Overhead



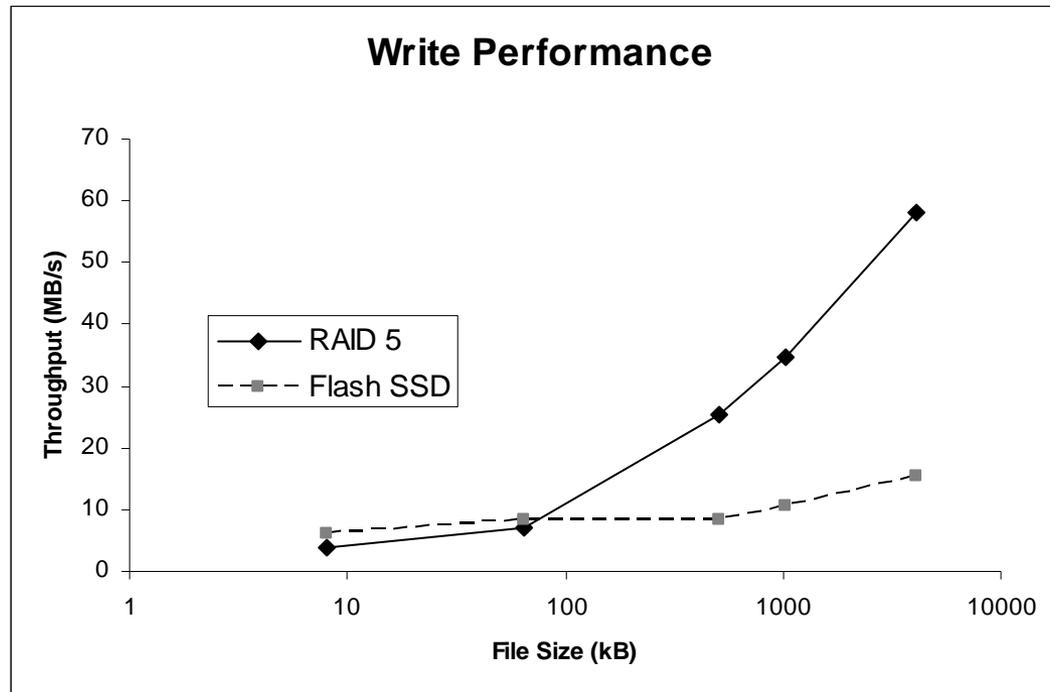
# I/O Limited Benchmarks



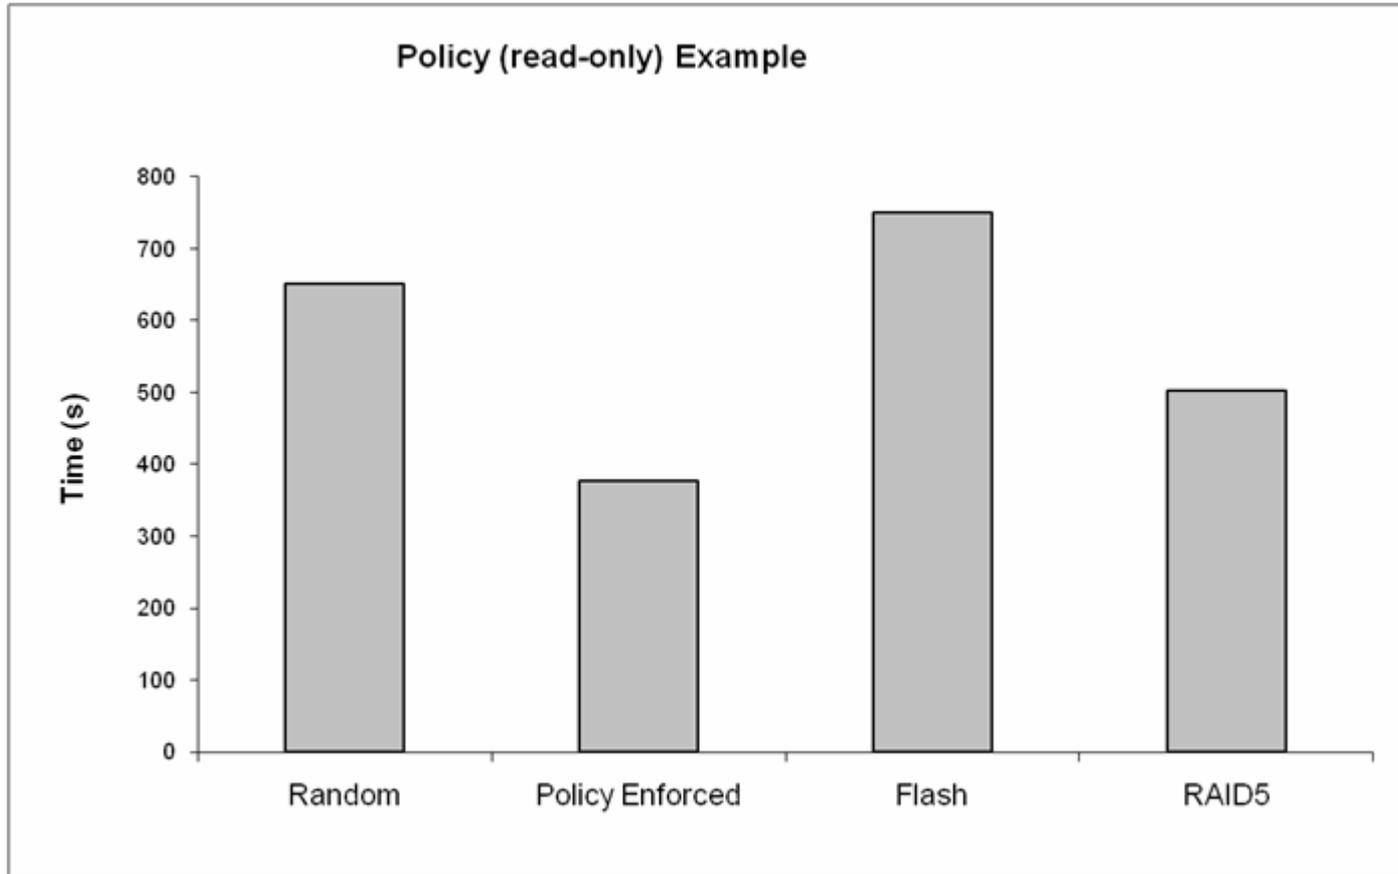
# Flash vs. RAID5 Read Performance



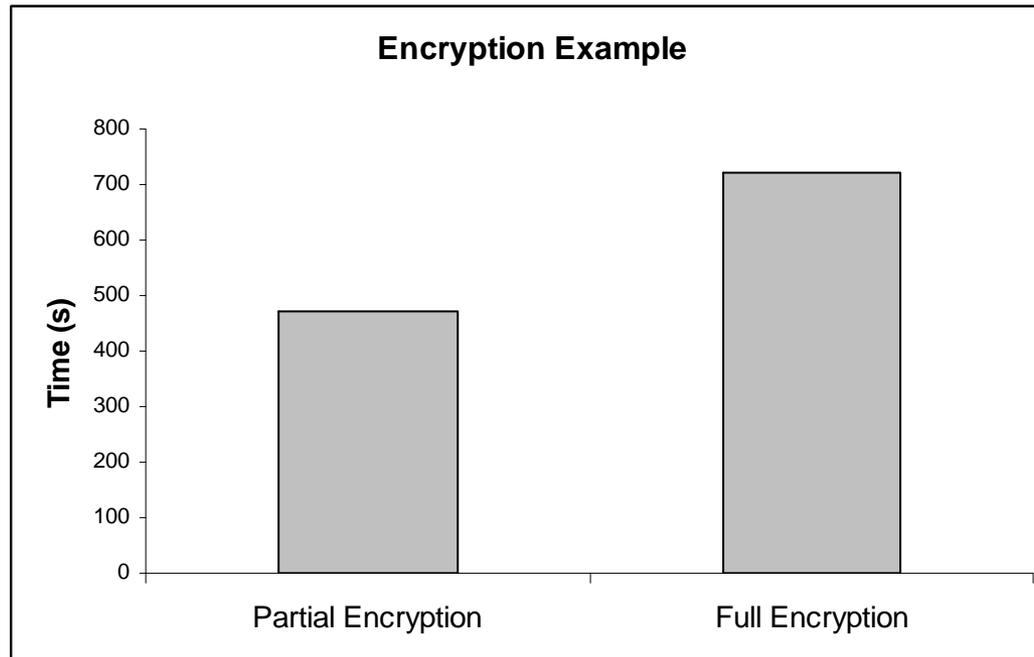
# Flash vs. RAID5 Write Performance



# Flash and Disk Hybrid System



# Disks with Encryption hardware



# Conclusion

- **Policy-based storage: Umbrella File System**
  - Allows matching application characteristics to devices
- **Little overhead for most scenarios**
- **Performance gain depends on employed policies**
- **Future work:**
  - Power management across heterogeneous devices
  - Automatic generation of policies

<http://cegroup.ece.tamu.edu/techpubs/2007/TAMU-ECE-2007-06.pdf>