

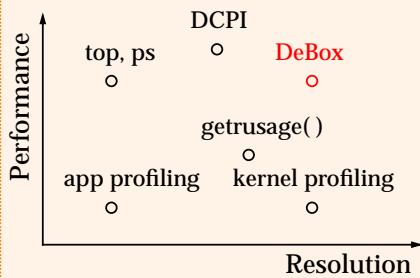
# Making the "Box" Transparent: System Call Performance as a First-class Result

## Performance Debugging of OS-Intensive Applications

Example: Web servers with SpecWeb99 workload

- High throughput CPU/Network
- Large dataset Disk activity
- Dynamic content Multi-process
- QoS requirements Low latency
- Workload scaling Low overhead

## Current Tradeoffs



## DeBox: Splitting measurement policy & mechanism

Cross-boundary debugging:

Add *primitives* into kernel

Programs control profiling

- Profile by call site and invocation
- Pick which processes to profile
- Selectively store/discard/utilize data
- Enable/disable, specify level of detail

In-band channel reporting:

Return feedback with each syscall

- *First-class result*, like *errno*

## DeBox Architecture & Implementation

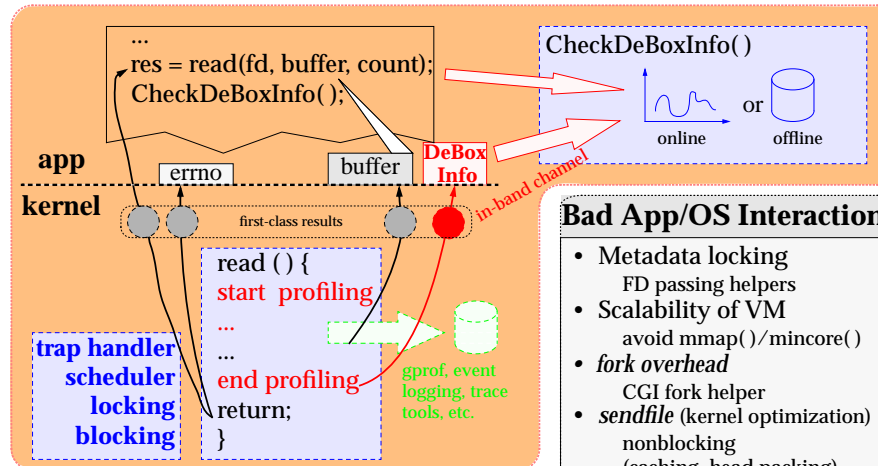


Table 1: SpecWeb99 Blocking Counts

resource label /counts	system call name /counts			
biord/166	open/162	read/3	unlink/1	
inode/127	readlink/84	open/28	read/9	stat/6
getblk/1	close/1			
sfpsys/1	sendfile/1			

Table 3: DeBox Overhead (time in usec)

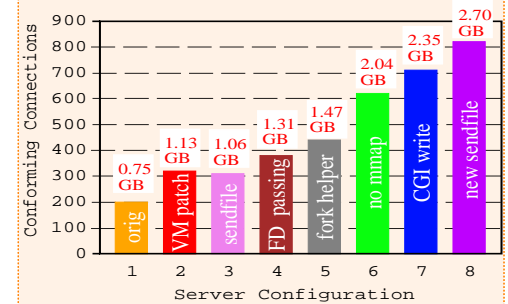
call name	base time	basic off	basic on	trace off	trace on	
						tar-gz a direcotry with
						kernel
getpid	0.46	+0.00	+0.50	+0.03	+1.45	
gettimeofday	5.07	+0.00	+0.43	+0.03	+1.52	
pread	256 Bytes	3.83	+0.00	+0.59	+0.26	+2.02
	1024 Bytes	6.74	+0.00	+0.68	+0.27	+2.02
	4096 Bytes	18.43	+0.03	+0.74	+0.29	+2.16
base time	275.61 msec		3078.50 msec		236.96 s	
basic on		+0.97 msec		+22.73 msec		+1.74 s
full support		+1.03 msec		+44.58 msec		+7.49 s

Table 2: Sample Output (copying a 10MB mapped file)

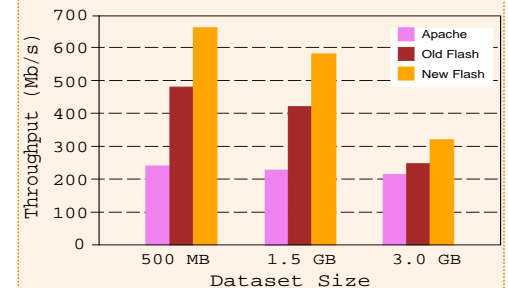
Basic DeBox Info	
4	system call #
3591064	call time (usec)
989	# of page faults
2	# of PerSleepInfo used
0	# of CallTrace used
PerSleepInfo[0]	
1270	# occurrences
723903	time blocked (usec)
biowr	resource label
kern/	file where blocked
vfs_bio.c	
2727	line where blocked
1	# processes on entry
0	# processes on exit
PerSleepInfo[1]	
325	# occurrences
2710256	time blocked (usec)
spread	resource label
.....	

## DeBox + Flash Server

### SpecWeb99



### Static Throughputs



### Diskbound Latency

