

# NFS Read/Write Performance

Mike Kupfer

*kupfer@Eng.Sun.COM*

# Some Benchmark Options

- LADDIS
  - highly configurable
  - aggregate performance
  - only measures server
- Modified Andrew Benchmark
  - measures client and server
  - mix of operations
  - no knobs, small working set
- Single process read/write
  - measures client and server
  - easy to configure
  - measures only a few operations

# Statistics

- *Art of Computer Systems Performance Analysis* (Jain, 1991)
- how many passes?
- keep track of standard deviation
- use more passes, higher confidence level when it's important
  - stripe size: easy to change, okay to be a little off
  - massive code change: important to be right
- these tables use 10 passes, which may be small

# Server Caching

- users probably want no-cache number
- shows impact of local filesystem

**Relative throughput when server cache is flushed between passes.**

	tcp		udp	
	read	write	read	write
UFS V2	-56%	-1%	-61%	-1%
UFS V3	-85%	-2%	-86%	-4%
SSA V2	-8%	-8%	-3%	-10%
SSA V3	-14%	-16%	-20%	-16%

# Version 2 vs Version 3

Version 3 throughput relative to Version 2.

	tcp		udp	
	read	write	read	write
SSA+RAM	88%	44%	68%	47%
SSA+RAM+Pr	83%	31%	58%	28%
UFS	-52%	95%	-45%	120%
UFS+Pr	-52%	11%	-45%	25%

Pr = Prestoserve

RAM = NVRAM (SSA fast writes)

# TCP vs UDP

**NFS/TCP throughput relative to NFS/UDP throughput.**

	Version 2		Version 3	
	read	write	read	write
SSA+RAM	-10%	-7%	0%	-9%
SSA+RAM+Pr	-13%	-9%	0%	-6%
UFS	1%	1%	-12%	-11%
UFS+Pr	0%	1%	-11%	-10%

Pr = Prestoserve

RAM = NVRAM (SSA fast writes)

# NVRAM (1)

- SPARC Storage Array 100 “fast writes” (4 MB NVRAM)

**Throughput improvement due to fast writes (no Prestoserve).**

	tcp		udp	
	read	write	read	write
Version 2	0%	245%	0%	261%
Version 3	0%	5%	0%	6%

# NVRAM (2)

- SBus Prestoserve card (4 MB (?) NVRAM)

**Throughput improvement due to Prestoserve (udp only).**

	Version 2		Version 3	
	read	write	read	write
SSA+RAM	0%	14%	0%	0%
UFS	0%	76%	0%	0%

SSA+RAM = SSA with fast writes