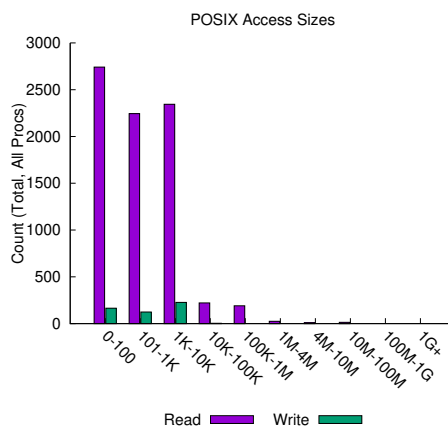



jobid: 123456	uid: 81434	nprocs: 1	runtime: 3793 seconds
---------------	------------	-----------	-----------------------

I/O performance *estimate* (at the POSIX layer): transferred **890.8 MiB** at **158.14 MiB/s**

I/O performance *estimate* (at the STDIO layer): transferred **33.3 MiB** at **97.55 MiB/s**




Most Common Access Sizes  
(POSIX or MPI-IO)

	access size	count
POSIX	8191	1142
	32	295
	4096	171
	4	132

File Count Summary  
(estimated by POSIX I/O access offsets)

type	number of files	avg. size	max size
total opened	648	1.2M	689M
read-only files	201	99K	11M
write-only files	3	48K	66K
read/write files	11	67M	689M
created files	14	53M	689M




## Average I/O per process (POSIX and STDIO)

	Cumulative time spent in I/O functions (seconds)	Amount of I/O (MB)
Independent reads	0.535551	922.865104675293
Independent writes	4.547624	1.25359916687012
Independent metadata	0.8914020000000003	N/A
Shared reads	0	0
Shared writes	0	0
Shared metadata	0	N/A

## Data Transfer Per Filesystem (POSIX and STDIO)

File System	Write		Read	
	MiB	Ratio	MiB	Ratio
/cvmfs/larsoft.opensciencegrid.org	0.36192	0.28871	783.77256	0.84928
UNKNOWN	0.17560	0.14008	0.00000	0.00000
/cvmfs/dune.opensciencegrid.org	0.65235	0.52038	139.06578	0.15069
/global/cscratch1	0.06373	0.05084	0.02677	0.00003



*sequential*: An I/O op issued at an offset greater than where the previous I/O op ended.  
*consecutive*: An I/O op issued at the offset immediately following the end of the previous I/O op.

#### Variance in Shared Files (POSIX and STDIO)

File Suffix	Processes	Fastest			Slowest			$\sigma$	
		Rank	Time	Bytes	Rank	Time	Bytes	Time	Bytes