

Symbolic IO™ Innovation Roundup

Rob Peglar
SVP & CTO
SYMBOLIC IO



The Sermon

- ★ All good sermons have three points and a joke
- ★ The topic is **The Brave New World of Big I/O**
- ★ So here it comes...ready or not...

Left to our Own Devices...

- Can you identify the device here? —————>
- Sadly, we're still doing I/O much like this
 - Blocks (512B, 4KB, etc.)
 - SSDs are no different
 - Address + offset (LBA)
 - The POSIX Five
 - read, write, open, close, ioctl
 - App wants to persist something
 - CPU gives up
 - Hands off the duty to the peripheral subsystem
 - RAM <—> DMA <—> Bus <—> Cache <—> Media
 - Context switches and interrupt handling
 - **Can we please stop doing this to ourselves? Please?**



What if We...

- Can you identify the machine here? —————>
- What if we went Back to the Future?
 - Persistent Memory
 - Block I/O to disk was the last resort
 - Memory was dear
 - Optimization ruled the day
 - The address space was wide enough
 - We didn't really need virtual memory (see prior bullet)
 - Programmers went to great lengths to write efficient code
 - We used self-modifying code - AI by definition (ML)
 - All of the above equalled the fastest computing on the planet
 - Bar none
 - **"In-Memory Computing" defined:**
 - **If your data isn't in memory...you're not computing**



Redemption at Last

- Can you identify the machine here? —————>
- The Advent of Byte-Addressable Storage
 - Persistent Memory (StorModules)
 - Block I/O to disk is not done (translated to compute, load/store bytes)
 - Memory is cheap - if you know how to use it correctly
 - Never persist raw data in memory - wasteful, inefficient
 - Persist the vectorized symbolic representation of it - compute results
 - Optimization rules the day
 - The address space is wide enough (48 bits physical = 256 TB per socket)
 - We don't really block storage (see prior bullet) except for 2nd tier/archive
 - Systems go to great lengths to persist efficient data
 - DNA-like encoding - markers, not raw data - amplify memory by 10X or more
 - We use self-modifying data - AI by definition (ML)
 - All of the above equals the fastest computing on the planet for its size & \$
 - We are now always computing in-memory by definition
 - **Because the best I/O is no I/O - so Big I/O becomes Big Compute**





Wasn't That Fun?

Rob Peglar
SVP & CTO
SYMBOLIC IO

