



Back to the Future: NAS as a Data Architecture

Brian Wong
Distinguished Engineer
Sun Microsystems



We've Been Here Before

- NAS has been around for 20+ years
- Actually, we're still here
- NAS hasn't changed that much
 - Files
 - POSIX
 - Handful of protocols
 - Protocol evolution has been measured
 - Vastly better ease-of-use

The Changing Processing World

- New processing architectures
 - More horizontal
 - More object-oriented
 - More distributed
 - More ubiquitous
- Are data persistence architectures keeping pace?

The Changing Application World

- First: data was about transactions about world artifacts
 - Banking, inventories, telco billing
- Now: data often is the artifact
 - Desktop publishing, especially digital media
- Next: data is about states of artifacts

The Changing Data World

- More zeros
 - More TB (or PB)
 - More files
- More, different types of access rights
 - Ownership no longer follows legal boundaries
 - Everything online all the time
- Keep that data!

Some Things Never Change...

- Data is still the core of processing
- Cost of maintenance is always too high
 - Always too complex to manage data
- Data is almost never near all the processing sites
 - If it is, there are too many copies
- Recovery is ever more important

Integration

- Solutions exist for many of these problems
 - Sometimes more than one
 - Some are pretty good, others aren't
 - Some are mature, others distinctly aren't
- Most users will need more than one
 - Most of them don't work with others
- Integration will be crucial

NAS Architecture

- Shares semantics of data
- Fine-grained: files, not devices
- Strongly heterogeneous
- Inherently virtualizes storage
 - NAS operations operate on data, not devices
 - Simpler for humans – and code
 - Most NAS architectures virtualize away devices

Blocks vs. NAS

- Can block storage address these data problems? Of course!
- NAS can solve them more effectively
 - Shared data semantics allow storage to be higher function, more versatile
 - Deployed solution architectures have better “division of labor”
 - Easier to offer heterogeneous access without resorting to lowest common

Technology Keeps Moving

- Massive application of VLSI
- Cost and value realignment of fundamental components
 - Memory – Disk – Tape
 - CPU – Network – Storage
 - Networks – Backplanes – Interconnects
- Software capabilities continue to grow
 - Particularly important: common

Performance

- Three dimensions
 - Throughput (MB/ sec)
 - Latency (elapsed ms/ operation)
 - Efficiency (μ s of CPU per operation)
- NAS latency already at par with blocks
- Two components of throughput
 - Pipe width – 1 Gbit NAS vs. 2 Gbit FC
 - Parallelization (pNFS and other efforts)

Efficiency

- The NAS performance issue has been efficiency – particularly on the client
 - TCP/ IP is vastly less efficient than SCSI
 - More semantics to implement with less hardware
 - TCP/ IP efficiency was 10- 20x lower than SCSI
 - Gap is now ~20% (tolerable) for key apps

Technological Convergence

- Modern NAS has clear value in ease-of-use and resultant TCO
- Improved security removes a major barrier to adoption
- Efficiency unlocks NAS values
 - NAS will be compelling when VLSI enables TCP/IP efficiency on par with SCSI/FC
 - Particularly true if HW encryption is included

Concurrent Transitions

- There are two separate transitions:
 - Transport changes from dedicated storage (FC) to commodity (IP)
 - iSCSI adoption is making IP-based storage much more common
 - Forces consideration of security for IP
 - Semantic changes from blocks to NAS
- Transport changes will tend to enable or accelerate the

Architectural Realignment

- Disk, tape and memory economics will cause persistence hierarchy changes
 - Much more memory, slightly less disk and tape
 - Make no mistake, we will still have tape
 - Probably some new ways of using tape too
 - NAS will be the way the hierarchy is addressed
- Combining NAS semantics and

APIs

- New or changed APIs driven by:
 - New application architectures
 - Higher function persistence systems
 - Changes in programming models
- Data persistence may well be embedded into new runtimes
 - Java runtimes
 - .Net classes

What Does This Mean?

- VLSI networking will overcome the primary obstacle to NAS deployment
- Security will get considerable attention
- The architectural shift from blocks to NAS will enable more sophisticated data persistence solutions



Back to the Future: NAS as a Data Architecture

brian.wong@sun.com

