Something about NFS, maybe

Brian Pawlowski
Vice President and Chief Architect
beepy@netapp.com

Some topics

- NFS today
- It’s the application
- Linux clusters
- A pointless aside
- iSCSI
- Doubling back to NFS
- Community

NFS Today

- "It was 20 years ago today."
  - SCSI and NFS grew up together
- Transformed from something you turn on in a UNIX release to a well-defined storage segment
- Home directories
- Large partitionable tasks that may run as parallel threads
  - Typical applications include search engines, e-mail, animation and rendering, scientific simulations, and engineering
- Scalable databases
- GRID computing

The Old Way

Imagine Charlton Heston in a chariot.

The New Way

Imagine an airplane full of chickens.

Scalable compute cluster

- Linux is ahead of the game
  - growing infrastructure, expertise and support
- It's all about choice!
- No! It’s all about freedom!
- Well, no actually, it’s all about cost.

Compute cluster points

- The x86 platform won
  - Any questions?
- Support costs may still be significant
  - ...but largely offset by the hardware cost savings - it's about leveraging small MP commodity x86 hardware
  - Some customers choose to pay more for better quality in order to lower support costs and improve performance
  - Maturation of "free software" - paying for support
- For Unix environments, NFS is the cluster file sharing protocol of choice
  - Customers simply want storage solutions that scale as easily as their compute clusters
  - But things change…

A pointless aside

Modern numerology

86	The preferred architecture for commodity components
2	Number of physical processors in commodity pizza boxes (poor man's blade)
128	Maximum expected nodes in a Linux database cluster
2000	Typical number of Linux nodes in a render or ECAD simulation farm today
10000	Expected number of nodes in Linux compute cluster in next two years?
<9	Number of filers per 1000 Linux nodes in GRID
1	There is only one - Linus Torvalds
10?	The number of trusted minions to Linus

Back to our regularly scheduled program…

Putting iSCSI in its place

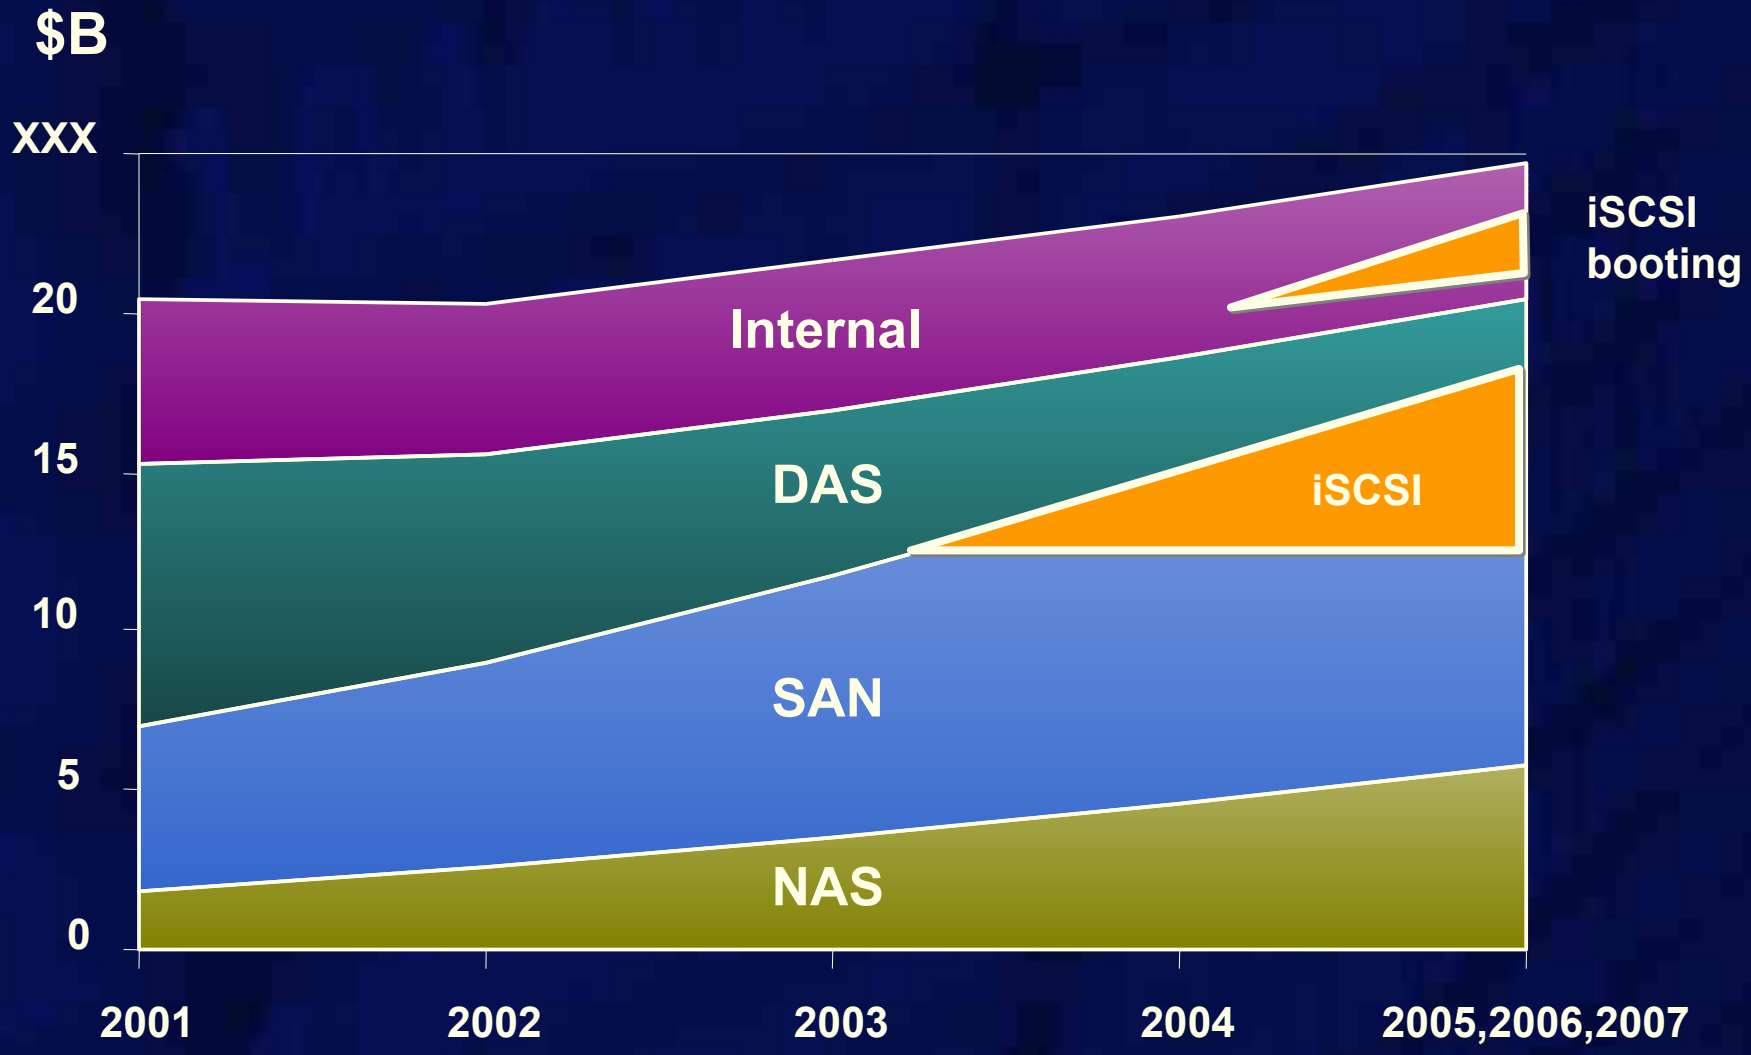

NetworkAppliance®

“So, iSCSI is a replacement for NFS, right?”

- In the first iSCSI presentation I made to a prospect, this was first thing out IT manager’s mouth
- I used to say “No”, but the thoughts underlying the question are interesting

iSCSI value proposition

- Leverage existing Gigabit → 10 Gigabit networking infrastructure
- Leverage existing rich set of management tools
- Leverage existing base of skilled personnel

But wait…

This sounds all so familiar

The NAS market redux

- Leverages traditional networking
  - CIFS and NFS run on TCP/IP and Ethernet
- Defines a model for sharing and collaboration
  - Scalable architecture
- Lower TCO

iSCSI points

- iSCSI software drivers freely and ubiquitously available
  - Windows platforms
  - Linux, and other *ixes
- HBAs and TOEs
  - Able to scale performance from software solution, to HW assist to full offload
- Saying “performance” and “iSCSI” in the same breath though misses the point
  - Performance is not always the primary issue
  - Many application deployments have spare (CPU and I/O) capacity
  - Optimize performance as needed

My sole contribution to NetApp marketing

Freely licensing iSCSI

iSCSI represents the path of least resistance

- It is semantically equivalent to FC SAN (SCSI)
  - But more familiar because of TCP/IP and Ethernet - so friendly outside the data center
- Application migration is trivial
  - My remote booting desktop from FC to iSCSI
- Provides a path for easily reclaiming FC port capacity by moving less critical apps to iSCSI
- With some of the important cost benefits of NAS

beepy, this is an NFS conference…

It's about applications

- Applications drive storage choices
  - What does the application vendor support?
  - What do they recommend?
  - For example, Exchange is driving iSCSI in the Windows environment
- Mix of applications in a single enterprise
  - There is no one perfect storage approach
  - There's likely more than one vendor

It's about data management

- Integration of applications with data management
  - Key applications like Exchange - application-driven backup/restore
  - Fertile ground for virtualization - blurring line between client application and storage
- Disaster recovery
- Finding data when you need it
  - Higher level data organization and grouping?

It's about cost

- Ability to (re)provision, expand and manage storage to maintain high utilization will most affect overall cost long term
- Leveraging commodity networking
  - iSCSI and NFS are similar here
- Primary storage and Nearline support for all storage access - transparently
  - Migration and replication
- Consolidation to reduce management costs

Let's put this in perspective

- "Wow. Michaelangelo, great statue - was that a 7 inch chisel you used?"
- "Great flick Welles, what camera did you use?"
- "Great quarter you guys had! Did you use NFS to access your financial data?"

Understanding the context around NFS

- That other operating system drives fundamental architecture decisions
  - Blade provisioning via NFS is a non-starter perhaps - because of multi-OS support
  - Enter iSCSI - the least common denominator
- People don't buy NFS servers
  - They buy Oracle applications
  - They build application compute clusters
  - And manage the data around it - with NFS perhaps

beepy, are you saying there is no difference between storage architectures?

Differences are important

- NAS protocols define a file view - higher level organization and semantics
  - Enables sharing
  - Enables large compute clusters (>5,000 nodes)
- iSCSI, like FC SANs, provides simpler SCSI block interface
  - Higher level semantics via explicit file system encapsulation
  - Sharing via layered cluster file system (complexity and cost?)
- Customers will use and continue to explore a variety of approaches

NFS Version 4 in one slide

- Driving conversation in NFS
- Emerging production releases
  - Customers are getting worried
- Fixing long standing problems
  - Security
  - Reliability
  - Performance
- Basis for future innovation
  - RDMA
  - pNFS
  - But what about migration/replication???
  - Name space???
  - Management?

Other than that Mrs. Lincoln…

- NFS = Network File System
- NFS = Not For Speed
- NFS = Not For Security

But you may sit there and think “My side of the boat is dry!”

Exactly.

Community needs deeper collaboration

- No surprises
  - Customers really want a better NFS Version 3
  - Are we prepared to provide support for NFS Version 4?
- Reliability
  - Testing
  - Scalability
- Playing well with others
  - Agreeing on common administration models
  - Agreeing on common features (else we will drop things from spec in IETF)
- Security
  - Administration needs to be simplified simplified simplified
- Performance is at bottom of list I think

Questions