



The New (Old) Challenges for Technology Infrastructure Today

Greg Bishop

Technical Advisor, Systems Operations
America Online, Inc.

8 April 2003

IEEE/NASA Mass Storage Systems and
Technologies Conference

What this talk is NOT



NOT!

What this talk IS

1. Some boring ancient history
2. Some depressing recent history
3. A trite description of where we are now
4. A slightly more interesting view of the road ahead



Ancient History...

(in computing)

- Computers were very expensive
 - Disk was outrageously expensive
 - Memory was precious
 - CPU costs of \$1/min were not uncommon
- People costs were significantly lower than hardware costs
- Result: companies spent money on people to manage technology extremely tightly.

(But you already know this...)

The 'Old' Challenge...

- Minimize hardware costs by maximizing hardware utilization (and thus save the company money).



The Technology Boom...

- Hardware costs became reasonable
- Qualified technical people were expensive and scarce
- Companies unrealistically believed that the money poured into technology would revolutionize business dramatically improve profitability
- Cost and utilization were not an object. This was as an arms race.
- Result: Companies spent money on technology with little regard to business cases.

(Yeah, Yeah...)

The Challenge of the Boom...

- Use the latest technologies and scale, scale, scale.



The Technology Bust...

- Hardware costs continued to drop
- Companies realized that they had overextended in technology expenses
- Qualified technology people became more available
- Companies began to expect a return on investment in technology (again).
- Result: Corporate spending on technology dropped significantly

(Yesterday's news...)

The Situation Today...

- Hardware costs continue to drop quickly – especially on servers (with a migration to commodity hardware).
- Software is becoming significantly more complex, requiring both more hardware and more people to manage it.
- Qualified technicians are available, though still somewhat expensive.
- A demonstrated return on investment is now required for technology.
- Result: Companies are now managing technology teams trying to lower costs in both hardware and personnel.

The 'New' Challenge...

- Lower total cost of technology. This includes people, capital and other costs (such as maintenance and license fees).
- Increase system capacity (especially in amount of required storage space).
- Improve scalability and reliability of systems.

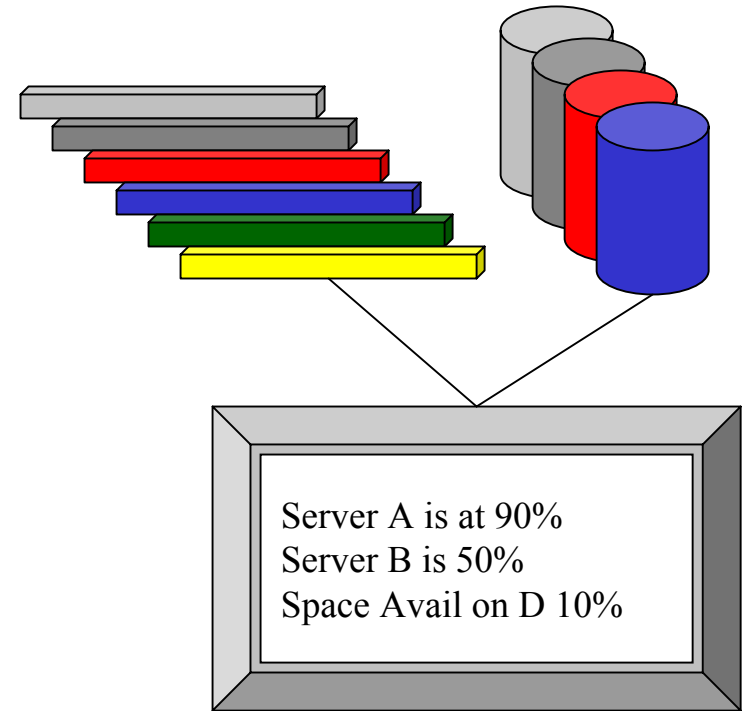


Lower Cost of Technology ...

- Improve system utilization
- Minimize personnel requirements for better management of hardware
- Migrate to lower capital cost technologies where possible
- Use the hardware and software that is right for the job

Manageability: A Major Need Today

- Any sustainable management platform needs to be based on standards adhered to by all
- GUI interfaces do not scale, interface needs to be a combination of GUI, database and CLI interfaces
- Any management platform needs to provide a consistent view of utilization across all vendors
- A single vendor solution is unacceptable



Manageability Summary...

- We need tools that easily manage hardware from a wide range of vendors using standards that all hardware and management platform providers abide by

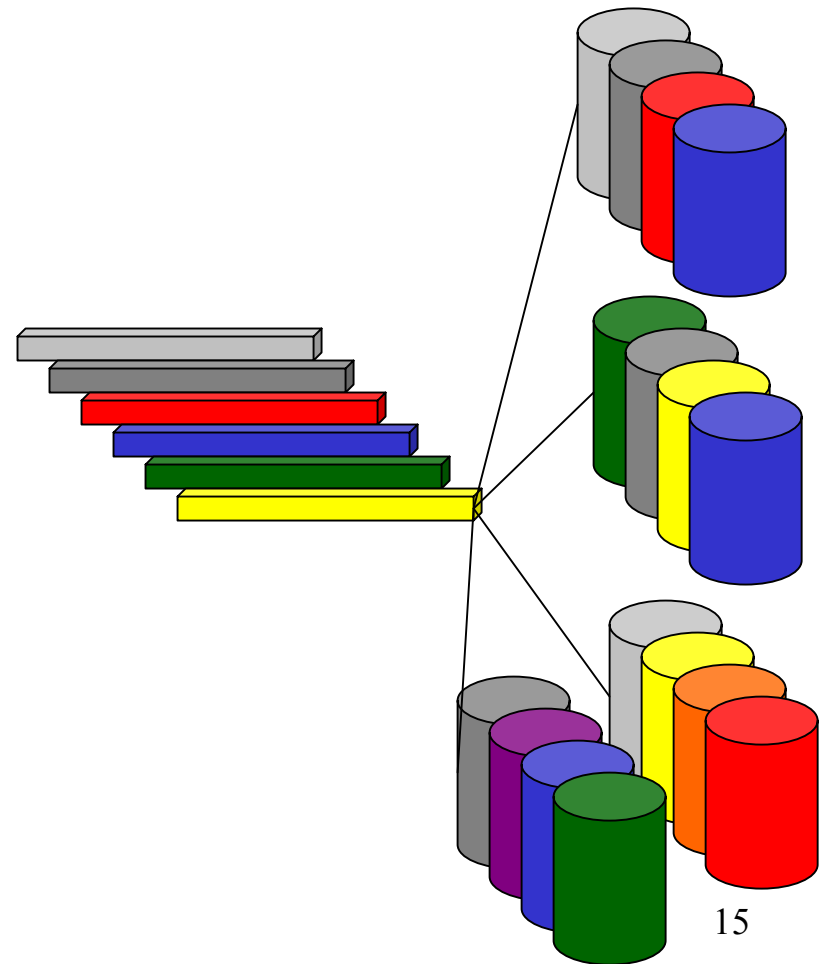


What is likely down the road...

- Dramatic increases in storage requirements
- Expectations of enterprise level disk costs will drop to levels similar to commodity disks (as is going on in the server market today)
- Expected reliability levels similar to those of the telephone company
- Maintain performance levels
- A need for cheaper transports. (Today you can spend more on the FC cards for a commodity server than the rest of the server!)

Storage Requirements...

- Marketing applications are requiring more detailed historical data on usage patterns...
- With more consumers moving to broadband, they are doing more with service providers (thus generating more transactions)
- Service providers want to become a central storage utility for customer data



Storage Costs...

- Executive management is currently expecting to see a huge savings by moving to commodity servers
- A number of disk system vendors have announced disk systems using commodity disks (though the pricing is still well above what can be purchased through retail outlets)
- Executive management is going to expect huge savings with use of commodity disks



Reliability...

- Current disk systems are quite reliable, but increasing reliability dramatically increases cost
- The increasing density of disk drives makes recovery times significantly longer when a failure does occur
- We need to find ways to increase reliability and improve recovery time without dramatically increasing cost



Performance...

- With the increasing amount of data per disk, the throughput per amount of data on the disk is dropping
- During recovery times, performance is significantly degraded today
- More spindles to improve performance again increases impact of failure (and frequency on a per piece of data basis)



Better transport...

- The move to lower cost hardware is in direct competition with the current use of FC for disk access
- We need to find a better way to transport the data from the disks to the new commodity servers
- There are currently two main approaches to solving this: NAS and iSCSI
- Most work appears toward NAS at this time



iSCSI

- Risky today. Not an entrenched standard
- Will require significant investment in new routers and switches for most enterprises
- Does not currently solve many of the problems faced by enterprises
- To work well in a commodity market, we need high performance/multi-vendor distributed file systems
- Has potential



NAS

- As it exists today: NAS=BAD
- More specifically: NFS=BAD
- CIFS does not scale economically
- Most development today is centered on making better servers (more scalable, better reliability)
- This work does not improve the situation...



What's Wrong with NFS?

- NFS servers are actually fairly good today (at least some of them)
- NFS protocol is not too bad...
- NFS clients are horrible and make NFS nearly unusable in most medium to large scale environments and completely unusable in production environments
- What needs to be done?



What can be done?

- Fix NFS clients
 - Realistically, this will most likely require that the protocol be changed
 - Client vendors **NEED** to write optimized drivers for their platforms
- Make a new Network file access protocol standard built for today's needs
 - Scalability
 - Reliability
 - Security
 - Performance
 - Support for all OS's



How can this happen?

- Unlike what you might think by talking with a number of technologists and CTOs, executive management does not care how these needs are met
- Most technologists and CTOs have pet technologies that they think will solve these problems and many of them will resist many other possible technologies, as they do not expect much time before company executives are going to expect the migration to commodity disk systems

In Short...

- Companies are going to need significant amounts of enterprise class storage with dramatically lower capital expenditure over the next couple of years
- Manageability and Performance are both needed to help companies meet usage goals (and thus cost goals)
- A significant number of technical challenges need to be faced as part of creating the solutions to the challenges ahead




Wild Guesses for the Future...

- In the not too distant future, non-rotating (not disks) storage systems will become viable as mass storage devices (TB of data level)
- The entire paradigm for storing data will need to change with a change to quantum computing (a huge revolution for the compute intensive environment)
- Network transport speeds will increase faster than Moore's law over the next 3-5 years

Questions?



End of Presentation



Appendix

California or Herring Gull (from conference Web site)



Royal Tern talking to Sandwich Tern



Osprey



Wood Stork



Willet



Ring Billed Gull



Green Heron



Royal Tern



Great Black Backed Gull



Laughing Gull



Great Blue Heron



Sandwich Tern



Semipalmated Sandpiper



Marbled Godwit



Fiddler Crab



Heermann's Gull

