



PRESENTS

**NETWORLD INTEROP**

an INTEROP event

# Evolving File Systems for Network Storage

---

Craig S. Hughes

Manager, Storage Software Development  
Network Engines, Inc.

September 11, 2001

# The Battle Between Performance & Management

---

- Trying to balance capacity, manageability and performance in storage systems is *not* easy.
- Will we have to trade highest performance for simplicity of management?
- The answer for now is probably “yes” (although performance profiling and simple tuning can get you a fair amount), but in the longer term, intelligent, adaptable software will help us get the most from our current and future systems.

# The Dilemma....

---

- General purpose configurations don't provide the best performance, and there are a bewildering number of configuration variables that need to be adjusted. This is true of DAS, NAS and SAN.
- The first step in getting our hands around these problems is by characterizing the applications accessing the data

# Example Applications

---

- **Web serving** – very many small files. Quick directory access and low per-file delivery overhead are key.
- **Streaming Media** – in most cases, very big files that must be delivered in uninterrupted streams. For these applications, file system or disk caches typically just ‘get in the way’, and disk data layout becomes very important.
- **Databases** - different types of access (highly localized access to indices and broad access to table data for example); may require file system or disk cache partitioning to provide best performance.
- **Backup** – ability to provide snapshots or very quick transfers of data are essential. Data is probably going to be accessed in large amounts one time only – caching doesn’t help that much except as a staging buffer for the disks.

# The Bad News...

---

- There aren't a lot of industry-wide 'best practices' for storage management that can be easily transferred to an IT staff. It takes experience - on your nickel
- Sophisticated administrators want access to all these 'knobs and dials' to optimize their performance...it's still a black art (good for job security, bad for those who can't afford the storage 'artisan'!)

# The Good News

---

- You *can* do a lot with general-purpose file systems in the short term, especially for NAS, where other OS features are bottlenecks. Novel caching mechanisms and meta data operation optimization (journaling/soft updates) have significantly increased performance of current file systems, and don't require complex tuning.
- Developing configuration profiles for classes of applications ease administrative headaches and provide basic tuning.
- In the long run, architectural overhauls will be needed: directory layout, data layout; user/application-defined caching policies and clustering will all help performance scale. Combined with adaptive, self-tuning monitoring sub-systems, file systems will not be the number one performance target on system administrators lists.