How to Give a Bad Talk

John Ousterhout, Tom Anderson, Dave Patterson, ...
(Channeled by Mike Dahlin)

I. XXX handwritten slide: thou shalt not be neat XXX

-- why waste research time preparing slides

Ignore spelling grammar and legibility

Who cares what 50 people think?

II. Thou Shalt Not Illustrate

Clients	xFS BW	NFS BW
1	5.71995e+05	1.65997e+06
8	4.425325e+06	1.19731e+06
16	1.095445e+07	7.88792e+05
32	1.38927e+07	4.70548e+05

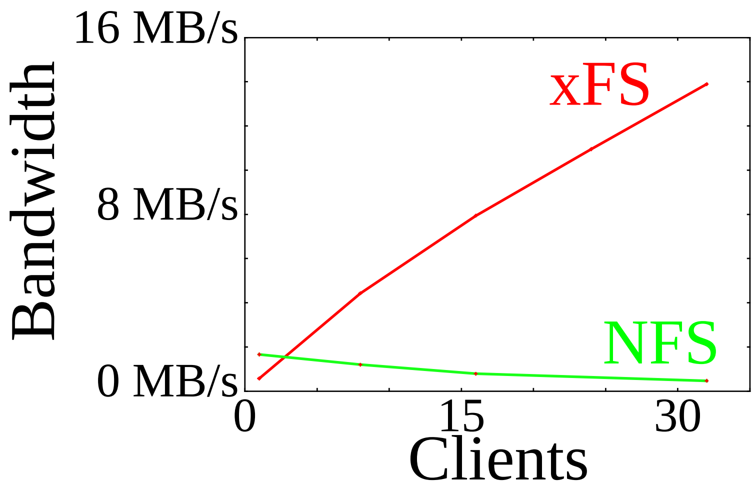

Table:

- Precision
- Allow Audience to Draw on Conclusions

Pictures:

- Confucious: “Picture = 10K Words”
- Dijkstra: “Pictures are a crutch for weak minds”
  Who are you going to believe?

III. Thou Shalt Not Covet Brevity
IV. Thou Shalt Not Print Large

Figure 8 illustrates the performance of our disk write throughput test, in which each client writes a large (10 MB), private file and then invokes sync() to force the data to disk (some of the data stay in NVRAM in the case of NFS and AFS.) A single xFS client is limited to 0.6 MB/s, about one-third of the 1.7 MB/s throughput of a single NFS client; this difference is largely due to the extra kernel crossings and associated data copies in the user-level xFS implementation as well as high network protocol overheads. A single AFS client achieves a bandwidth of 0.7 MB/s, limited by AFS's kernel crossings and overhead of writing data to both the local disk cache and the server disk. As we increase the number of clients, NFS's and AFS's throughputs increase only modestly until the single, central server disk bottlenecks both systems. The xFS configuration, in contrast, scales up to a peak bandwidth of 13.9 MB/s for 32 clients, and it appears that if we had more clients available for our experiments, they could achieve even more bandwidth from the 32 xFS storage servers and managers.

Figure 9 illustrates the performance of xFS and NFS for large reads from disk. For this test, each machine flushes its cache and then sequentially reads a per-client 10 MB file. Again, a single NFS or AFS client outperforms a single xFS client. One NFS client can read at 2.8 MB/s, and an AFS client can read at 1.0 MB/s, while the current xFS implementation limits one xFS client to 0.9 MB/s. As is the case for writes, xFS exhibits good scalability; 32 clients achieve a read throughput of 13.8 MB/s. In contrast, two clients saturate NFS at a peak throughput of 3.1 MB/s and 12 clients saturate AFS's central server disk at 1.9 MB/s.

While Figure 9 shows disk read performance when data are not cached, all three file systems achieve much better scalability when clients can read data from their caches to avoid interacting with the server. All three systems allow clients to cache data in local memory, providing scalable bandwidths of 20 MB/s to 30 MB/s per client when clients access working sets of a few tens of megabytes. Furthermore, AFS provides a larger, though slower, local disk cache at each client that provides scalable disk-read bandwidth for workloads whose working sets do not fit in memory; our 32-node AFS cluster can achieve an aggregate disk bandwidth of nearly 40 MB/s for such workloads. This aggregate disk bandwidth is significantly larger than xFS's maximum disk bandwidth for two reasons. First, as noted above, xFS is largely untuned, and we expect the gap to shrink in the future. Second, xFS transfers most of the data over the network, while AFS's cache accesses are local. Thus, there will be some workloads for which AFS's disk caches achieves a higher aggregate disk-read bandwidth than xFS's network storage. xFS's network striping, however, provides better write performance and will, in the future, provide better read performance for individual clients via striping. Additionally, once we have ported cooperative caching to a faster network protocol, accessing remote memory will be much faster than going to local disk, and thus the clients' large, aggregate memory cache will further reduce the potential benefit from local disk caching.

V. Thou shalt not waste space: transparencies are expensive!!!

VI. Thou Shalt Not Use Color

- Detracts from Seriousness
- Too Flashy
  Your Ideas Should Sell Themselves
- Emphasizes Some Words Over Others
- Encourages Illustrations
- Costly to Print

VII. Thou Shalt Not Skip Slides in a Long Talk

- You Did the Work to Prepare Them
- Better Approach: Go Faster
  - Shows Respect for Audience:
    - Not Trying to Slip Something Past Them
- Skip Summary and Conclusions if Necessary

VII. Thou shalt not make Eye Contact

- You should avert eyes to show respect
- Don’t acknowledge questions from audience

IX. Thou Shalt Cover Thy Naked Slides

- Keep Audience on Your Point
- Surprise Them With Your Train of Thought
  How else can they know how clever you are?
- Advanced Topics:

  Double Cover Slides

  = Double Advantage

  Stand Near Projector
  Several Advanced Techniques

X. Thou Shalt Not Practice

- Wastes Several Hours out of 2 Years of Research
- Spontaneity
- If You do Practice
  - Argue with any suggestion
  - Make talk longer than allotted time
- Most Important Commandment!

XXX Advanced topics -- combinations of techniques