


---

# Distributed File System Performance


**Milind Saraph / Rich Sudlow**  
**Office of Information Technologies**  
**University of Notre Dame**



# Questions to answer:

---

- **Why can't you locate an AFS file server in my lab to improve my performance?**
- **What kind of performance improvement can I expect if I use a 100 Mb/s network interface?**
- **What are the performance differences of shared vs switched ethernet?**
- **What are the performance differences of various RAID levels? How should disks be striped for the best performance?**
- **Are all RAID arrays equivalent?**
- **If and when we move/migrate to DCE/DFS, what are the performance implications?**




# Questions to answer: (cont.)

---

- **What are the effects of the server CPU speed on my client performance? Is it worthwhile to purchase MP server machines?**
- **How many AFS/NFS translators are needed to serve Mac/ PC clients? How are the translators best configured?**

**Our tests are an attempt at obtaining data in order to answer some of these questions.**



# Current Configuration:


---

- **AFS Servers:**
  - 9 Sun Ultra Enterprise Servers (5) 2/1170, (3) 2/2170, (1) 2/1200 with 128 MB RAM
  - 1 Sun SPARCserver 20/71 with 96 MB RAM
    - Approximately 500 GB disk space
      - 5 Servers have Sun Model 112, 63 GB SSA (RAID 0)
      - 2 have Sun Model 101 31.5 GB SSA (RAID 0)
      - 1 has a Sun Model 210 controller
        - 18 x 2.1 GB disks (RAID 0+1)
      - 1 has a MTI 9300 126 GB RAID Array (RAID 5)

**All AFS servers use a single 100 Mb/s network interface.**

---

**DECORUM '97**




# Current Configuration: (cont.)

---

- **AFS Translators (serving Mac and PC clients):**
  - 8 AFS/NFS Sun SPARCserver 20/71
    - 1 processor ( 75 MHZ SuperSPARC, 1 MB cache)
    - 512 MB RAM 2 x 1.05 Fast SCSI-2 disks
  - 1 AFS/AFP Sun SPARCserver 20/71
    - 1 processor ( 75 MHZ SuperSPARC, 1 MB cache)
    - 256 MB RAM 2 x 1.05 Fast SCSI-2 disks

**All AFS Translators use a single 100 Mb/s network interface.**



# Current Configuration: (cont.)


---

- **AFS Clients:**

- Approximately 400 Sun workstations
- 60 RS6K workstations, (16 node IBM SP1, 14 node IBM SP2)
- 30 SGI workstations, (12 node Origin 2000)
- Hundreds of PCs and Macs via translators. Most dorms networked. OIT provides software for AFS access.

- **AFS Volumes:**

- Approximately 10,000 active volumes; volumes range in size from 50 MB - 20 GB.



# Current Configuration: (cont.)

---


- **DCE Test Cell**

- A single server Sun SPARCserver 20 DCE/DFS test cell, migration plans in progress.

- **NFS Email spooler**

- A 8 processor Ultra Enterprise 4000 which acts as a general time share machine in addition to being the NFS mail spooler for approximately 220 Unix clients.

Note: Another 4 processor Ultra Enterprise 4000 also acts as a email server utilizing a AFS based server running POP.



# Test Configuration:

---

- ( NFS/AFS/DFS ) File Servers

- (1) Sun Ultra Enterprise 4000

- 4 Processor (167 MHZ, 512K cache), 512 MBRAM,  
Sun SPARCstorage Array Model 112 Array with 2.1 GB disks  
2 x 2.1 GB 7200 RPM FW SCSI-2 internal disks


- (1) Sun Ultra Enterprise 2/1170

- 1 Processor (167 MHZ, 512K cache), 128 MB RAM,  
2 x 2.1 GB 7200 RPM FW SCSI-2 internal disks  
Sun SPARCstorage Array Model 112 Array with 2.1 GB disks

- (1) Sun Ultra Enterprise 2/1200

- 1 Processor (200 MHZ, 1MB cache), 128 MB RAM,  
2 x 2.1 GB 7200 RPM FW SCSI-2 internal disks  
Sun SPARCstorage Array Model 112 Array with 2.1 GB disks

---



# Test Configuration: (cont.)


---

- **Clients**

- (6) Sun Ultra 1 Model 140 workstations
  - 1 processor (143 MHZ, 512 K cache), 64 MB RAM
  - 1.05 GB Fast SCSI-2 internal disk
  - 10/100 Mb/s Sbus Card

- **Network interface for clients and servers**

- David Systems 10 Mb/s shared Hub
- Cisco Catalyst 5000 switch
  - 10 Mb/s switched 12 port card
  - 100 Mb/s switched 12 port card



# Tests:


---

- **I/O Performance**

- Diskrate program from Convex (now HP).
- Flushed sequential reads/writes for 1, 4 and 16 and 64 MB file.
- RAID 0 stripes of 1, 2, 3, 4 and 6 disks on SSA Model 112 & single internal disk.
- Tests of UFS, NFS, AFS, AFS/NFS, DFS and DFS/NFS clients.

- **Network Performance**

- Netperf 2.1 from HP (UDP, TCP throughput and UDP, TCP transactions) varying socket sizes and number of clients.



# Results / Conclusions:

---

- **Test results can be found at:**
  - <http://www.nd.edu/~rich/Decorum97>
- **Realistic benchmarks of your cell are essential**
  - Simple reproducible tests allow you to develop a knowledge base for future capacity planning and problems.
  - The data which you obtain dispels “myths” and allows more intelligent and informed communication with management and customers.

