

# Distributed systems

- **Scale and performance in a distributed file system. Howard, J. H., Kazar, M. L., Menees, S. G., Nichols, D. A., Satyanarayanan, M., Sidebotham, R. N., and West, M. J. . ACM TOCS 6, 1 (Feb. 1988)**
  - **Distributed file system design for scalability: wide area and a large number of clients**
    - **Heterogeneous clients**
    - **Volumes and cell servers**



## AFS goals

- Location transparency: Same name space on all the machines
- User mobility: Users must be able to access their files from any machine
- Security: should not depend on the workstation
- Performance: good performance for desktop users
- Scalability: thousands of users?
- Availability: effects of failure restricted to failed components and not to the entire system
- Integrity: users should not feel the need for backups
- Heterogeneity: Different types of desktops



# Trust

- Workstations are not trusted. Communications between AFS and Workstation encrypted
- Vice (trusted component), Venus (runs on the client)
- Cache management: whole file cache, modifications are reflected in Vice on close() (directory modifications are passed through to Vice)
  - Assume cache entries are valid unless callback. On reboot, consider all items suspect (callback not delivered)
  - On open, check if cache available and no callback = use cached copy, else fetch copy from server
    - Already opened files continue to use local cache



## Consistency semantics

- Writes visible to all other processes in the same workstation. New opens will see changes on close
- Meta operation propagated immediately
- Volumes and volume movement (copy-on-write snapshot)
- Backups via volumes: read-only clone of volume

