# AIX Compute Farm

## IBM RISC System 6000 Cluster

## at

## The T. J. Watson Research Center


Paul Fink Jr.

Marc Donner    Robert Barrett    James Doherty    Michael Greenwood

Scott Kennedy    Edward Overly    Elizabeth Poole    William Pope

Gregory Rose    Angelo Rossi    Christine Sundaresan

IBM

*T. J. Watson Research Center*

*Yorktown Heights, New York*


May 2, 1991

# IBM T. J. Watson's Compute Farm

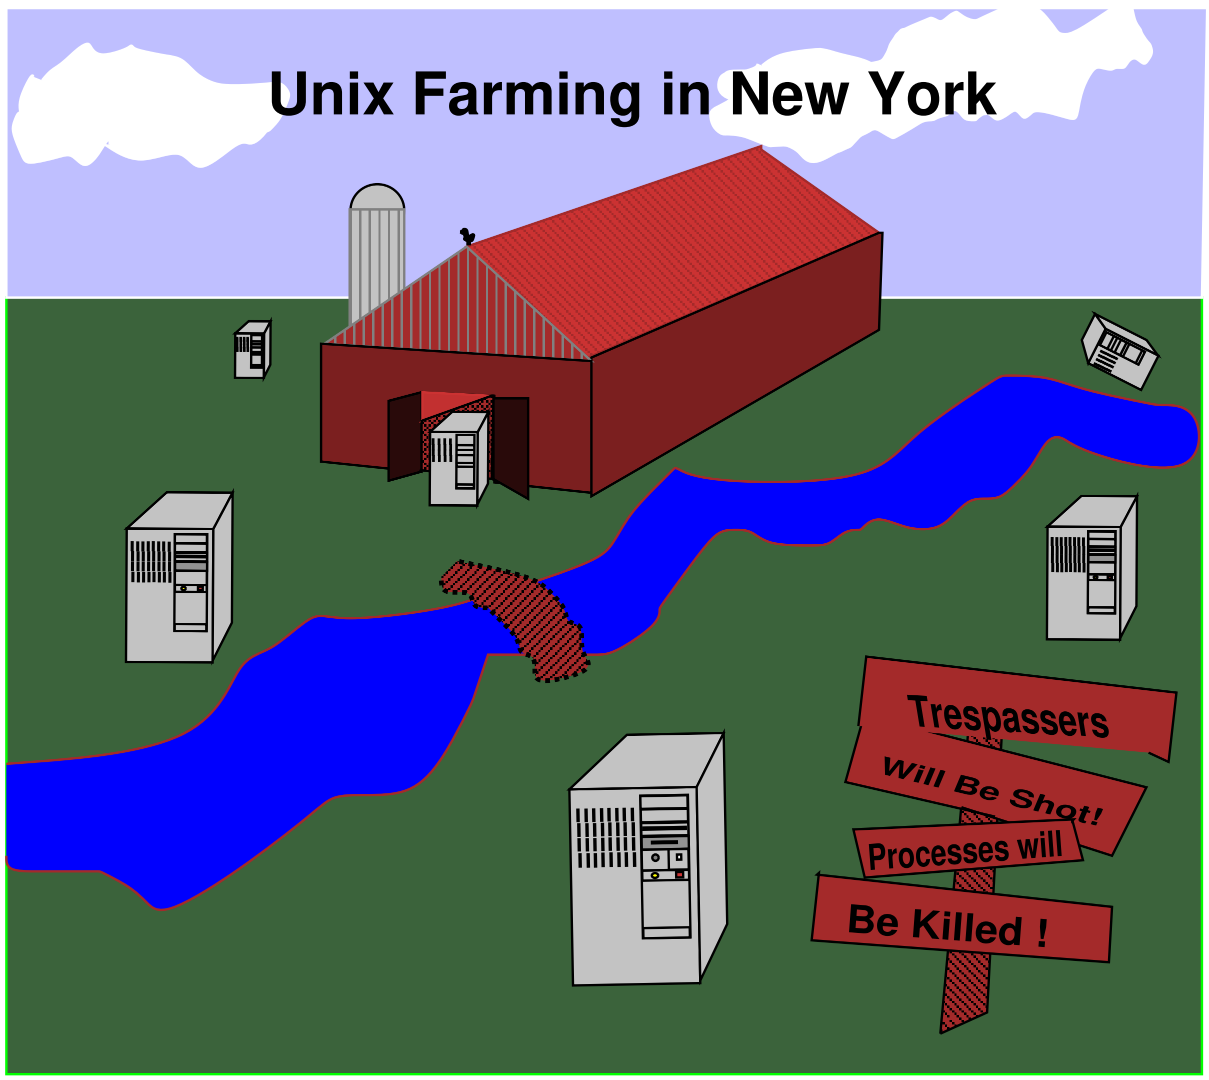

# Definitions

**Large** less than GB.

**Medium** approximately 100 MB.

**Small** tens of MB.

# Traditional Super Computing

- Computer centers are rated by how many MIPS that can provide.
- But our purpose is to advance science and technology.
- Prehaps a better metric is how well a computing center facilitates physical research.

# What would the ideal computing environment for the physical research be?

- Flexible & Extensible User Environment
  - Research is a creative activity and needs a creative environment.
  - Environment must allow for customization.
  - Facilitate collaboration with other researchers.
- Seamless Growth Path (3M to 3G)
  - Ability to move codes and data from desktop to supercomputer completely transparently.
  - Researchers never out grow available computing resources.
  - Put resources where they can be utilized the best.

- Program as Data
  - What's important is the results of program not the running of the program.
  - Researchers data should be available where it is needed without FTP or concerns about internal number representation.
  - Data should be available equally at the compute server, visualization facility, and on their desktop simultaneously.
  - Graphics and visualization techniques should be readily available and simple to use.
- Physics not CS
  - Researchers should be concerned with modeling and numerical techniques not programming languages and machine dependent optimizations.

# Why Cluster Workstations?

Unix networking tools allow clustering workstations for:

- **Greater peak performance**: single job parallelism
- **Economy of scale**: compute servers managed as a single machine.

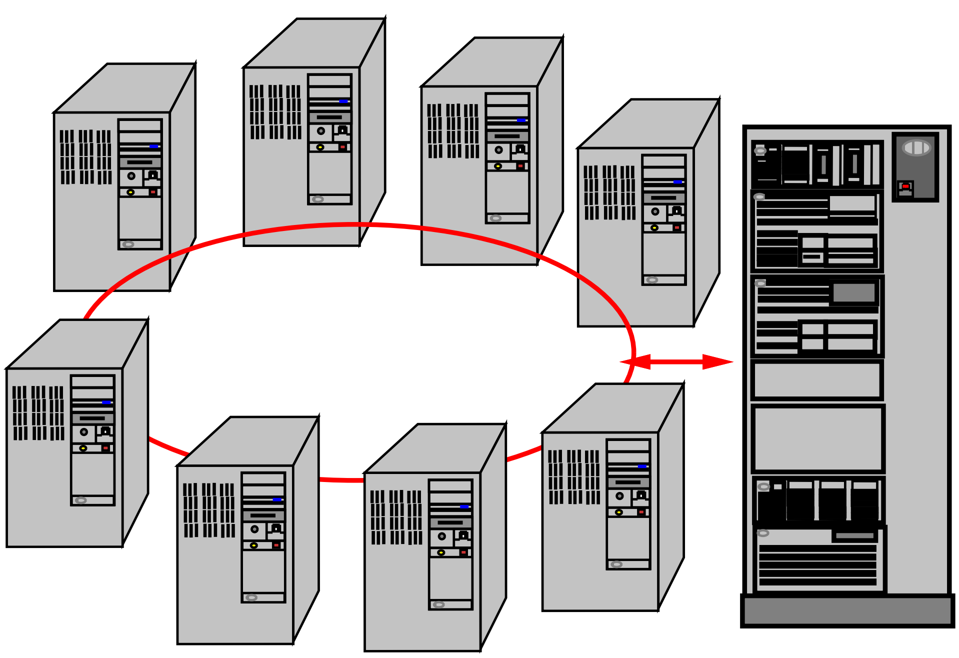

# Why T. J. Watson's Compute Farm?

- Hundreds of large scale users.
- High performance RISC workstations provide performance comparable with mainframes.
- However a single workstations can satisfy the needs of only a few simultaneous large scale users.

$\Rightarrow$ Cluster machines to obtain mainframe like multi-user performance and utilization.

# Goals of the Farm Project.

- A production system for use by T. J. Watson Researchers.
- Robust: uniting desktop workstations with central resources.
- Scalable: initially 15 compute servers.
- Flexible: incorporating new processor & network technologies.
- Use standard Unix utilities when possible.
- Operated and maintained by existing mainframe staff.

# Farm Architecture

- Homogenous, all AIX 3.1 but different speeds and sizes.
- Each machine can provide equivalent services.
- Appears as a single machine to users.
- Integrated with user's machines.

# T. J. Watson's Compute Farm

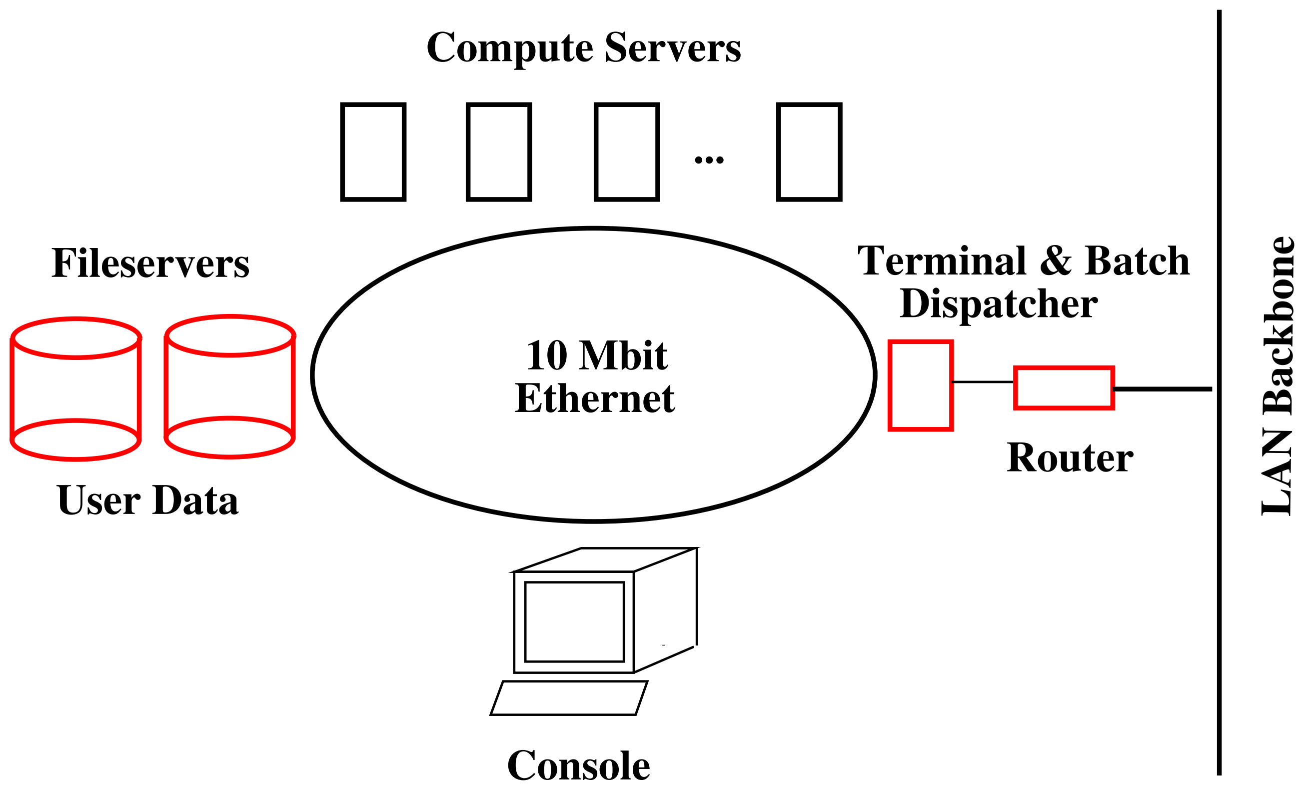

# Hardware

- Compute Servers
  - Model 540 with 256 MB RAM or Model 550 with 512 MB.
  - 2.5 GB of disk space.
    - $>$ 1 GB page space
    - 0.5 GB /tmp
    - No local user space

- Fileservers

  - Model 930 with 20 GB of disk space.

  - Holds all user data, `/u`.

- Networks

  - 10 Mb ethernet.

  - Dedicated router isolates Farm's network from building LAN.

  - Single IP address gateway for interactive logins.

  - Possible expansion to FDDI or dual ethernets.

# Farm Topology.

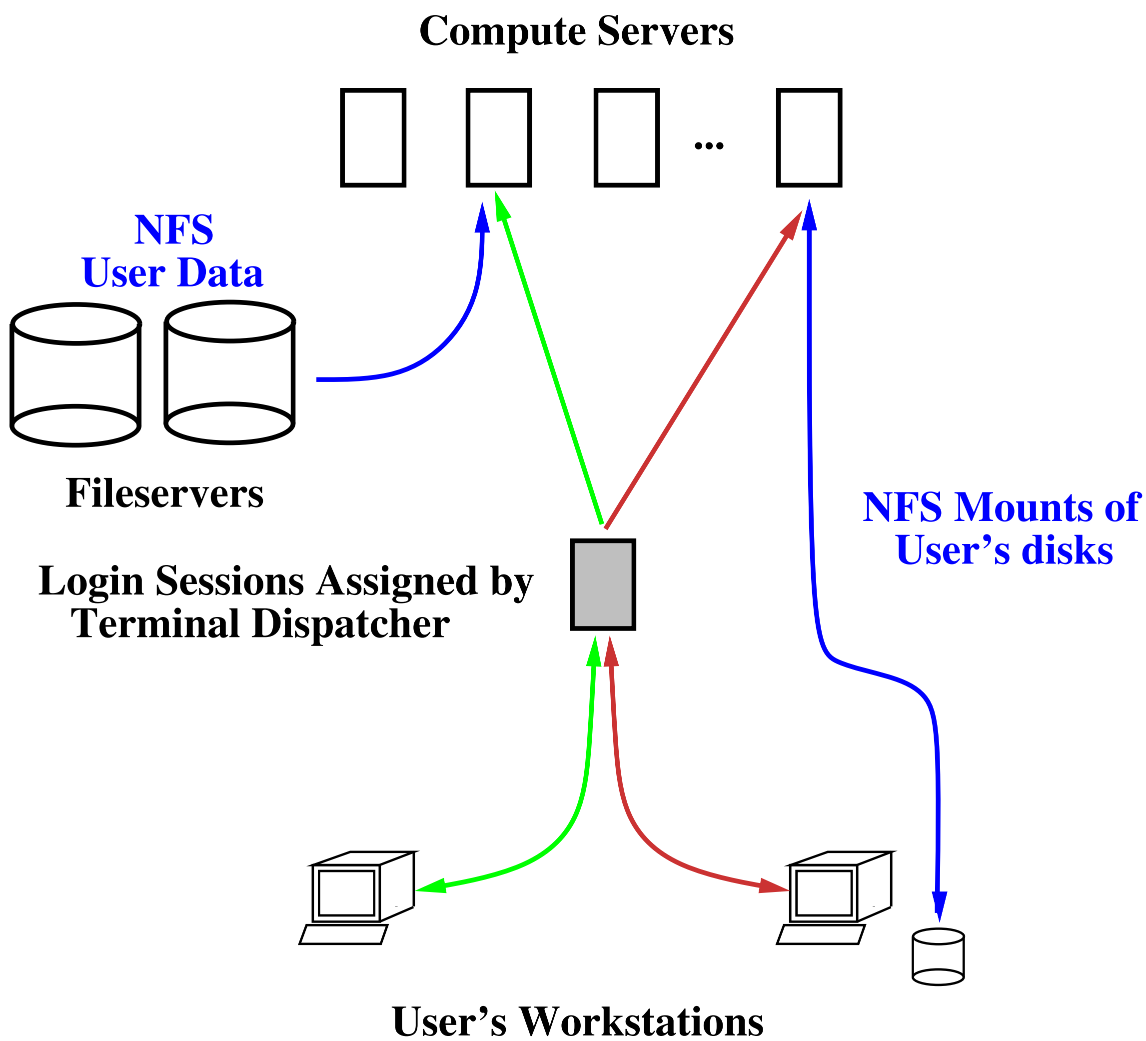

# Cluster as a Single Machine.

- IP Gateway
    - Users are given a single IP address for Cluster
    - Gateway routes logins to next available machine.
    - A user maybe on more than machine at a time.
    - Multiple logins place users on multiple machines.
    - IP address of cluster machines is available to users.

- Faking a unified process space

  - User commands are extended to work in a cluster environment.

  - Unified `PID` has two components: `host:PID`.

  - ps, kill, renice all operate on unified `PID`.

- Sample output from `ps -ef`.

| UID | PID | PPID | C | STIME | TTY | TIME | CMD |
|---|---|---|---|---|---|---|---|
| cliff | 7:15426 | 7:10817 | 0 | 11:35:27 | pts/0 | 0:00 | -ksh |
| twpoon | 7:17908 | 7:1 | 120 | 14:46:15 | pts/1 | 276:17 | mc15suov |
| pfink | 7:21221 | 7:14564 | 0 | 23:03:15 | pts/2 | 0:00 | -tcsh |
| cliff | 1:13716 | 1:11667 | 0 | 11:48:18 | pts/1 | 0:00 | -ksh |
| ggr | 1:19046 | 1:14181 | 0 | 15:58:22 | pts/0 | 0:00 | -ksh |
| cliff | 1:20746 | 1:13716 | 104 | 11:51:38 | pts/1 | 610:44 | rt madac |
| cliff | 2:10674 | 2:16713 | 54 | 17:30:42 | pts/1 | 157:27 | rt occ2 |
| twpoon | 2:13908 | 2:1 | 53 | 14:45:14 | pts/0 | 313:32 | mc15suov |
| cliff | 2:16713 | 2:20808 | 0 | 17:27:58 | pts/1 | 0:00 | -ksh |
| jimd | 4:14693 | 4:1 | 3 | 17:42:50 | pts/0 | 9:44 | xcpu |
| jimd | 4:17285 | 4:1 | 3 | 17:43:15 | pts/0 | 8:54 | xldave |

# Maintaining System Software.

- Single System Image.
  - All machines have identical system software & local tools.
  - There is one *System Image*.
- Rdist, **r**emote **dist**ribution, for updating.
  - Software is installed on a master and pushed out to slaves.
  - Hourly updates of user accounts, `/etc/passwd`.
- NFS used for user data.
  - All machines provide equivalent service to a user.

# Operations & System Administration.

- Managing Machines.
  - Cluster programs query a database of machine attributes.
  - Operators change the database to affect how the cluster is configured.
- Automated machine checkup program `DrDr`
  - Hourly checkup of system processes, disks, & networks.
  - Fault reports are sent to operators via same system as for mainframes.

- System Monitoring

  - Xgmon

    - Used to monitor site-wide network, workstations & Farm.
    - SNMP agents for CPU usage & load average.
    - DrDr reports integrated into display.

  - Standard Unix tools used by system administrators.

    - xps, xcpustate, xperfmon, xload.

# Integrating Farm with Desktop Workstations.

- Site-wide Single System Image.
  - Hierarchical level of updates.
- Single Home Directory.
  - Farm home directory may be NFS mounted on office machine.
  - Or home from office machine may be mounted on Farm.
- Using idle office machines: Condor.
  - Finds idle machines and adds them to batch pool.
  - Check-points job and removes it when user returns.

# What We've Learned, *so far*.

- Need for distributed filesystem.
  - `Rdist` is a pain.
  - Single server model not sufficient for large cluster.
  - More efficient model needed to utilize high speed fiber.
- Need more control than TCP/IP.
  - How do you reboot a large number of machines?

- Need for distributed queueing system.
  - Single system for site-wide batch submissions.
  - Allow scheduling of parallel jobs, e.g. linda, PMake.
  - Incorporate Condor paradigm.
- Unix workstations in a mainframe computer room.
  - Unix system administrator $\Leftrightarrow$ VM system programmer.
  - Unix systems are very complex and ill-defined.
  - Management tools for multiple machines don't exist.