The Limits That Await Us

Merritt E. Jones

The MITRE Corp, MC Z-250

1820 Dolley Madison Blvd

McLean VA 22102-3481

+1-703-883-5471; 703-883-5230 (FAX)

merritt@mitre.org

Presented at the THIC Meeting on April 23, 1997

DoubleTree Hotel, Falls Church VA 22043

MITRE

BACKGROUND

- In some fields, storage systems in the hundreds of terabytes to petabytes are routinely discussed and are being procured.
- Many of these systems are long range in outlook and will require a steady infusion of new technology
- This presentation looks (somewhat loosely) at areas of concern, forecasts, limits, and developments in storage related technologies

BACKGROUND (cont.)

- The distributed computing model as presented on the next chart is the model that most major vendors are implementing in their hardware and software
- Major components include computing, storage (data) management, networking, visualization, and systems management.
- Distributed storage systems have become an increasingly important part of the distributed computing model
- Distributed storage systems are typically comprised of multiple levels of storage including:
  - primary: processor CPU's, cache, memory
  - local: solid state disk, magnetic disk, optical disk
  - tertiary: tape drives, tape (and disk) robots
  - archival: shelf

Distributed Computing Environment

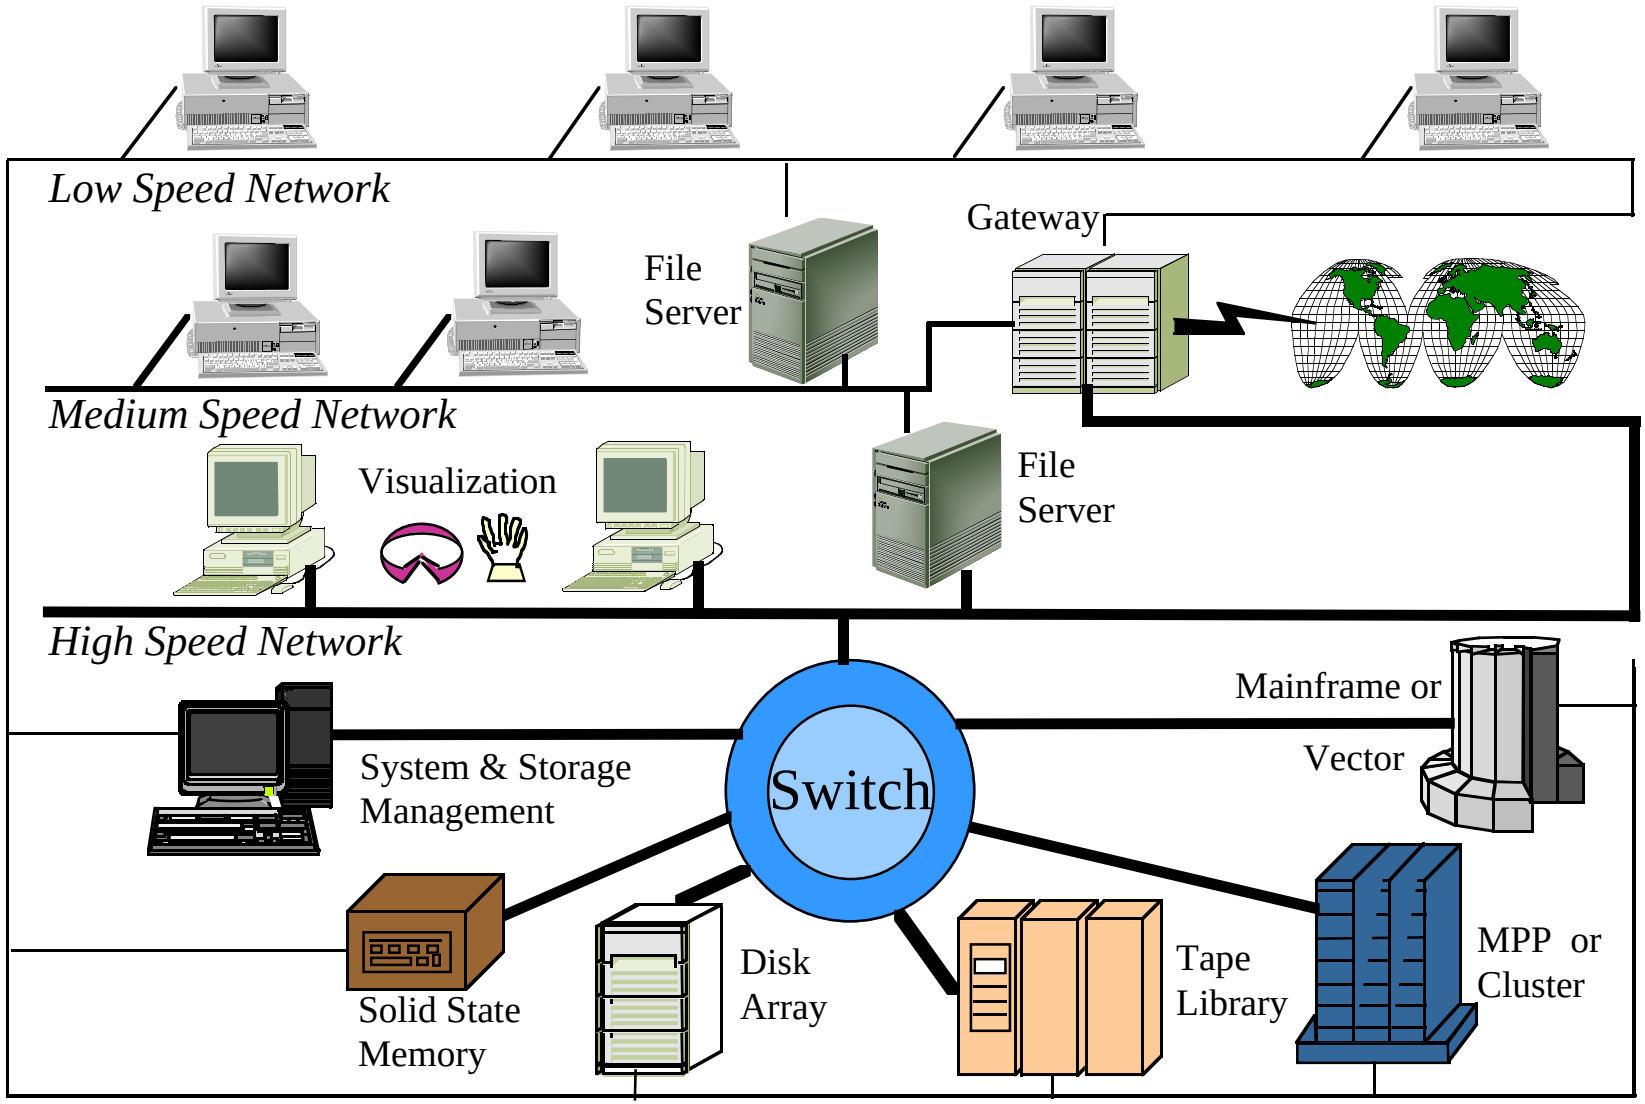

MITRE

Predictions - Something To Think About

- In 1901, two years before they took off from Kitty Hawk, Wilbur Wright told his brother Orville, that man would not fly for fifty years.
- “My figures coincide in setting 1950 as the year the world must go smash.” Henry Adams (American Historian), 1903
- The telephone system is in trouble because there won’t be enough women in the right age group to be switchboard operators. Circa 1955.
- "I have traveled the length and breadth of this country and talked with the best people, and I can assure you that data processing is a fad that won't last out the year." The editor in charge of business books for Prentice Hall, 1957
- It’s easy to predict the future because we’re always wrong. If you want to be right, read science fiction. Anon.

Predictions - A Little Closer To Home

- "Everything that can be invented has been invented." Charles H. Duell, Commissioner, U.S. Office of Patents, 1899.
- "I think there is a world market for maybe five computers." Thomas Watson, chairman of IBM, 1943
- “Computers in the future may weigh no more than 1.5 tons." Popular Mechanics, forecasting the relentless march of science, 1949
- "But what ... is it good for?" Engineer at the Advanced Computing Systems Division of IBM, 1968, commenting on the microchip.
- "There is no reason anyone would want a computer in their home." Ken Olson, president, chairman and founder of Digital Equipment Corp., 1977
- "640K ought to be enough for anybody." Bill Gates, 1981

A Few More Things To Think About

- The first disk drive, RAMAC, was introduced in 1956
  - Cost was $10,000 per megabyte. (Today, it is less than 25 cents per MB.)
  - Capacity was 5 MB and it was about the size of two refrigerators. It had 50 disks, 24 inches in diameter.
  - Areal density was 2 kbits/sq. inch compared to the best today of 1.32 gbits/sq. inch
- The computer that launched the Saturn V into orbit and sent the upper stages to the moon could perform 7.5K fixed point and 7.5K floating point operations per second
  - Today's luxury car has more computing capacity than the lunar module

And A Few More Things To Think About

- The (modest) PC (100 MHz Pentium) used to prepare this presentation has 16 times the memory, 60 times the processing processing power, and a bigger, faster disk than the $4M mainframe that supported the Mission Control Center for the Apollo Missions. (As seen in the movie, “Apollo 13.”)
  - The same modest PC has about 120 times the compute capacity of the original Shuttle Onboard Computers.
- A little mental calibration (by Roy Clickery @ CalTech)
  - 5 megabytes - the complete works of Shakespeare
  - 5.5 gigabytes - the lifetime reading for an average person
  - 10 terabytes - the Library of Congress printed material
  - 2 petabytes - all U.S. academic research libraries
  - 5 exabytes - all words ever spoken by human beings

REQUIREMENTS DRIVING DATA STORAGE TRENDS

- The following factors influence trends in density, seek times, read and write rates, and storage systems image
  - Multi-petabyte sized libraries
  - Multi-gigabyte sized files
  - Efficient access to increasing metadata
  - High speed database searches
  - Long distance file transfers
  - Very high rate data streams
  - Guaranteed response times
  - Technology induced transitions
  - Primary/secondary storage transparency
  - Cost
- It is worth noting that data requirements tend to linearly track increased CPU capacity

TECHNOLOGY GROWTH RATES

- There is considerable variability in different technologies and generalized trends must be recognized as such
- However, subject to some interpretation, major trends include
  - Processor power is doubling every 1 1/2 years
  - Memory density is doubling every 2 years
  - Magnetic media capacity is doubling every 2 years
  - Sustained I/O rates (single path) have gone from .5 - 3 megabytes/second to 2 - 20 MB/S in the last 10 years
  - Rotational delays have halved in the last ten years
  - Tape seek times are difficult to characterize due to thinner, longer tapes and multiple lengths for the same container

BASE TECHNOLOGY TRENDS - DENSITY

- The areal density improvement curve for magnetic media is continuing at 60% per year compounded. This should continue for at least the next 5 years.
- The highest demonstrated disk density today is 5 gigabits per square inch (3 times the best production density)
- Research is already at 10 gigabits per square inch and 25 megabytes/second
- The limit on the technology is currently believed to be about 70 gigabits per square inch (the superparamagnetic limit) and will be reached within the next decade
- There are about 10 million billion atoms per square inch on the surface of a disk (atomic surface density limit) attainable by a scanning tunneling microscope (already demonstrated.)
- An Atomic Force Microscope (under practical room conditions) has demonstrated 45 gigabits/sq. inch and could reach 300 gigabits/sq. inch.

BASE TECHNOLOGY TRENDS (CONT.)

- “Industry workhorse” tape thickness is 10 - 18 microns today. It will go to 7 - 8 microns. There are projects at 3-4 microns. Some opinion holds that 1 - 2 microns is feasible.
- The highest tape areal density today is 129 megabits per square inch (DDS-3)
- Volumetric density for magnetic tape should be 1 terabyte per cubic inch by 2000

FUNDAMENTAL TECHNOLOGY RULES

Here are some general rules of thumb that are relatively well known. They are collected in a recent article by John R Mashey (SGI) and Darryl Ramm (Permia, Inc.) in the September (96) issue of UNIX Review.

- ***Rule 1***. DRAM memory capacity increases by four times every three years.
- ***Rule 2.*** DRAM random-access time improves more slowly. Access times have only been cut in half in the last ten years.
- ***Rule 3.*** Microprocessor performance increases by about three to four times every three years, depending on benchmark and compiler choices.
- ***Rule 4.*** Small disks double in capacity every 18 to 24 months.
- ***Rule 5.*** Disks improve in access speed (seek time) by a factor of 1.3 every eight to ten years.
- ***Rule 6.*** Money can buy bandwidth, but latency is forever.

A Few Limits?

- The microprocessor era may be half over. James Meindl (RPI provost) suggests that chip miniaturization will end by the second or third decade of the 21st century.
- Robert Keyes, IBM Research Scientist, wrote (in 1993), "…when conducting areas approach to within 100 angstroms of one another, quantum effects threaten…" At that time lab researchers were at 30 angstroms. "…in commercial devices, maybe a decade remains before that limit is reached."
- An Intel extrapolation of Moore's Law (with a learning curve adjustment) predicts maximum computer processor capability of 10**7 MIPS. A speed of light limit for switching predicts a maximum speed of 10**11 MIPS.

A Few Limits? (cont.)

- Based upon switching elements needed to implement instructions, chip real estate vanishes between 2006 and 2013.
- Network capacity can sustain a tenfold increase every four years which exceeds Moore's law (both modified for the learning curve).
- The technical capacity of light-wave communciations (about 10,000 times today's networks) will exceed the ability of computers to load it around 2002.
- 2-3 terabits per second over long distances within the next decade.

Source: Collected by Ted Lewis, Naval Postgraduate School

High End Requirements - Some Examples

The following values do not represent a single program

- 600 GB/hour total input/output rate (1.33 gbits/second)
- 3 - 5 TB/day input stream
- 1 TB on-line disk caches (per server)
- 150 - 300 TB near-line storage
- Two known projects in the 5 to 10 petabyte total storage range
- 50 million files
  - 1 week old data in 5 - 10 seconds
  - two month old data in 30 seconds to 5 minutes
  - 10 year old data in minutes to hours
- 140,000 “active” cartridges
- 2 million total pieces of media
- thousands of simultaneous users
- 2000 data base queries/minute

Comments on Automated Media Exchangers

- Current exchange (one insertion and one ejection) rates are up to the 350 exchanges per hour range for large robots, 600 exchanges per hour for small robots. For most applications, this is not the limiting factor
- This will increase to about 450 exchanges per hour and stay there for a few years
- Most vendors will be introducing smaller, entry level robots
- The current upper limits on groupings and cartridges for a single systems image seems adequate
- More robots will handle a mix of tape technology
- Media residency time is increasing (due to increased capacity). Robots will tend to have more drives

Selected Technology Projects

0 For information purposes some technologies being developed/ investigated are listed below

- Optical tape drive by LOTS Technology, Inc. Goals: 1 terabyte capacity. 3490 form factor. 15 (initially) to 100 megabytes/second
- High performance, variable data rate, multimedia magnetic linear tape recorder by 3M, Seagate, et al. Goals: 30 -100 megabytes/second, 1/10 the cost, and performance exceeding high end helical scan systems
- Ultrahigh-capacity multilayer rewritable optical disk by Eastman-Kodak, et al. Goals: 1 terabyte of data on a 14 inch optical disk
- Electron trapping optical memory by Optex Communications Corporation. Rewritable 5.25 Inch optical disk. Goals: 4 gigabytes per surface, 55 - 250 megabits/second

SELECTED TECHNOLOGY PROJECTS
(CONT.)

- Holographic volumetric storage. Prototype goals: 100 megabits/second, 20 megabytes. Theoretical goals: 1 terabyte/second, 1 terabyte/cu. cm.
- Two-Photon based optical memory. 125 gigabytes/cu. cm.
- Additional technologies include
  - ion milling (Los Alamos National Lab). 600 gigabytes on a 5.25 inch platter
  - DNA (multi-wavelength) optical storage devices . 100 bits/ micron (about 100 times today's densities.)
  - photosensitive protein (bacteriorhodopsin) in halobacterium halobium. Hundreds of megabytes in milligrams of "ferment your own" storage.