



Slouching Towards Better Access Models: The Future of the NCAR Mass Storage System

George Williams

National Center for Atmospheric Research

1850 Table Mesa Drive, Boulder CO 80307-3000

Phone: + 1-303-497-1843 FAX: + 1-303-497-1848

E-mail: gwilliam@ncar.ucar.edu

Presented at the THIC Meeting at the Embassy Suites Hotel
Denver South

Englewood CO 80112

on June 28, 2000

The Premier Advanced Recording Technology Forum

THIC Inc.

Outline



- Background
- The NCAR MSS Today
- Challenge: Growth
- Strategies
- Challenge: Value
- Strategies
- Summary



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

1



Background: NCAR

■ National Center for Atmospheric Research

- Plans, organizes, and conducts atmospheric and related research programs
- Established in 1960
- Provides state-of-the-art research tools and facilities to the atmospheric sciences community
- Sponsored by the National Science Foundation



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

2



Background: SCD

■ Scientific Computing Division

- Provides, applies, and advances high performance computing technologies and expertise
- Terascale computing
- Networking
- Scientific visualization
- Data analysis
- **Mass Storage**



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

3



Background: The NCAR Mass Storage System

- In production since 1986
- Based on the *IEEE Storage Reference Model*, out of which grew the IEEE P1244 standard effort
- First large scale MSS to implement 3rd party transfers using a dedicated SAN (HPDF)
- Designed as an **archive**




Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

4



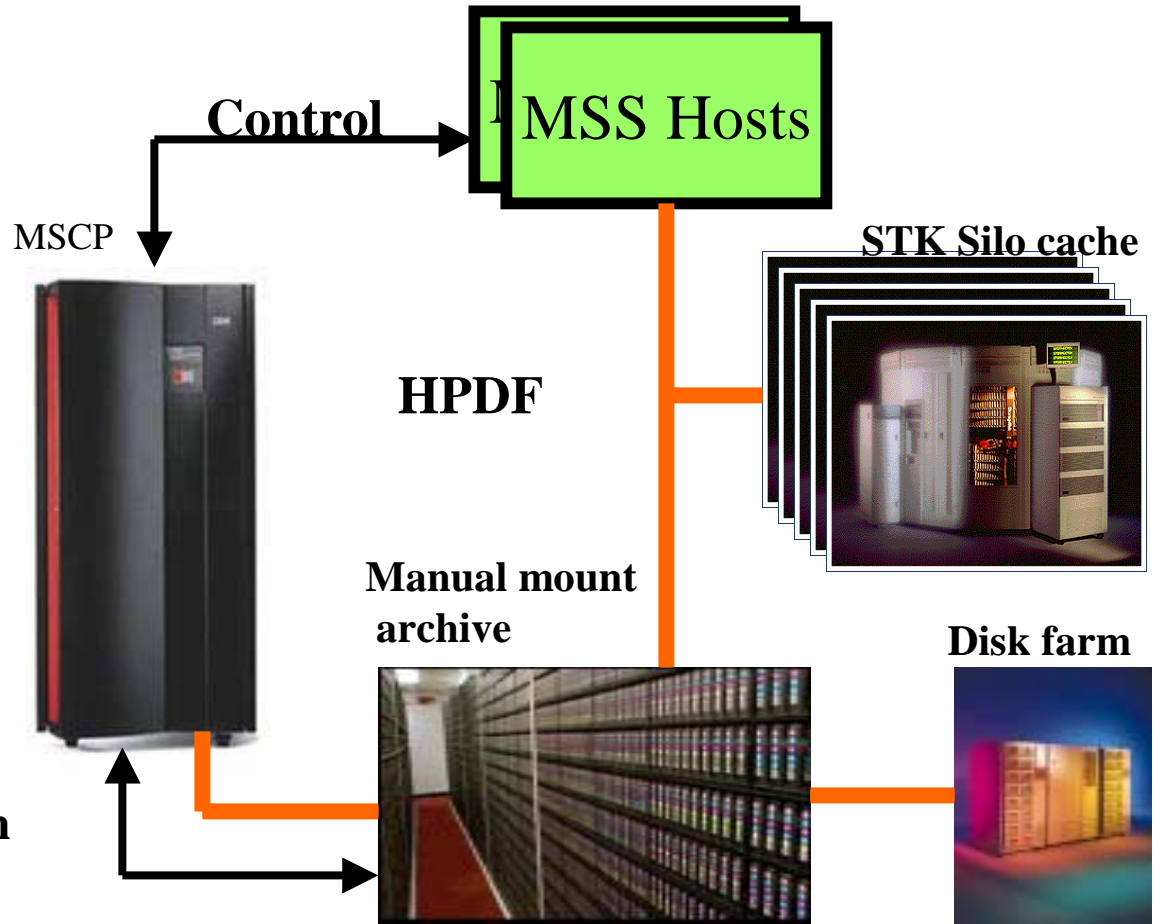
Background: Hardware

- 1986 - 3480
- 1987 - 90 GB Disk farm
- 1989 - 4480 Silo - 2.5 TB
- 1991 - 3490E
- 1993 - 4490 Silo - 5 TB
- 1993 - HIPPI
- 1994 - 180 GB Disk farm
- 1996 - SD-3 Silo - 300 TB
- 1996 - 3390 Disk farm
- 1999 - 9840 Silo - 120 TB
- 2000 - 9840 Silos - 240 TB
- 2000 - (9840 follow-on)



The NCAR MSS Today: Numbers

1986 - 3480
5 Silos
160,000 3490E cartridges
7,500 SD-3 cartridges
7,000 9840 cartridges
 1993 - 4490 Silo - 5 TB
255 TB data
8,000,000 Files
 1996 - 6 TB/month growth
50 TB/month moved
19 MB/sec sustained
 1999 - 9840 Silo - 120 TB
500,000 Transactions/month



The NCAR MSS Today: Miscellaneous

- Designed as an archive, used as a file server
- Growth proportional to available compute power (114 Bytes/MFLOP)
- Easy-to-use but non-standard interface
 - UNIX-like file names
 - rcp-like commands
 - Extended “user metadata” (limited)




Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

7




Challenge: Growth




Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000



**Scientific
Computing
Division**




Growth: Projected GFLOPs



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000



**Scientific
Computing
Division**



Strategies

- Traditional approaches will help...
 - New tape technologies
 - More silos
 - New channel/network technologies
- ...but projections paint a gloomy picture
 - MSS budget would have to **increase** significantly
 - Computing budget would have to **decrease**...



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

10



Strategies

- MUST limit growth
 - Change user behavior
 - Provide proper tools
 - Implement “Classes of Service”
 - Adjust charging



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

11



Challenge: Value

Traditional user needs

1. Store my data
2. Move it quickly
3. Do it cheaply
4. Let me access it from multiple servers
5. By the way, here's some more data
6. Don't bother me

Additional user needs

7. By the way, I have these computers, too
8. WhyCantIDoThat.com
9. By the way, here's some more data
10. And tell me again why I need you?

"Things fall apart; the centre cannot hold..."



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

12



Strategies II

- Global Shared File Service
 - A set of centrally-managed high-performance SAN-enabled shared file servers
 - POSIX file system interface
- Catalog/Search Service
 - A flexible web-accessible database that stores scientific metadata, similar to a search engine
 - Higher-level, content-based interface

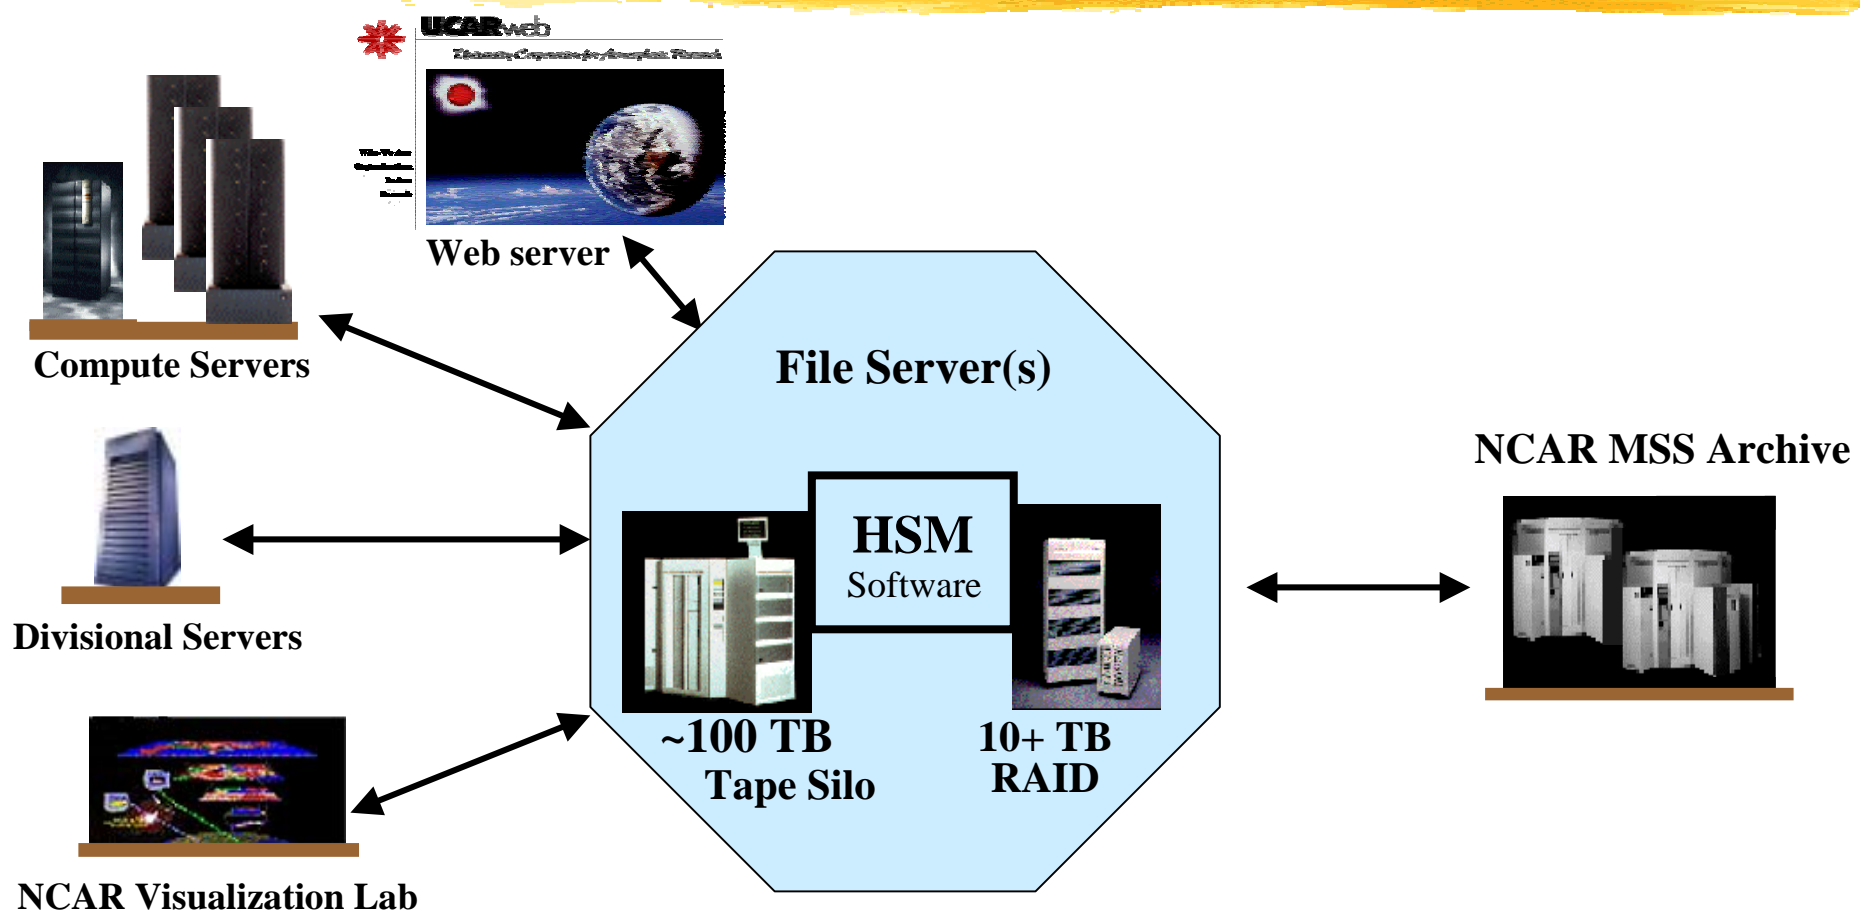


Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000


13



Strategies II: Global Shared File Service




Traditional Unix NFS



Data pipe to client is usually slow, NFS server is a bottleneck



Vendor - Partitioned Disk (current day SAN solution)



Sharing a file requires making a duplicate copy; additional disk space is required, it takes time to make the copy to be made, and coherency is a problem

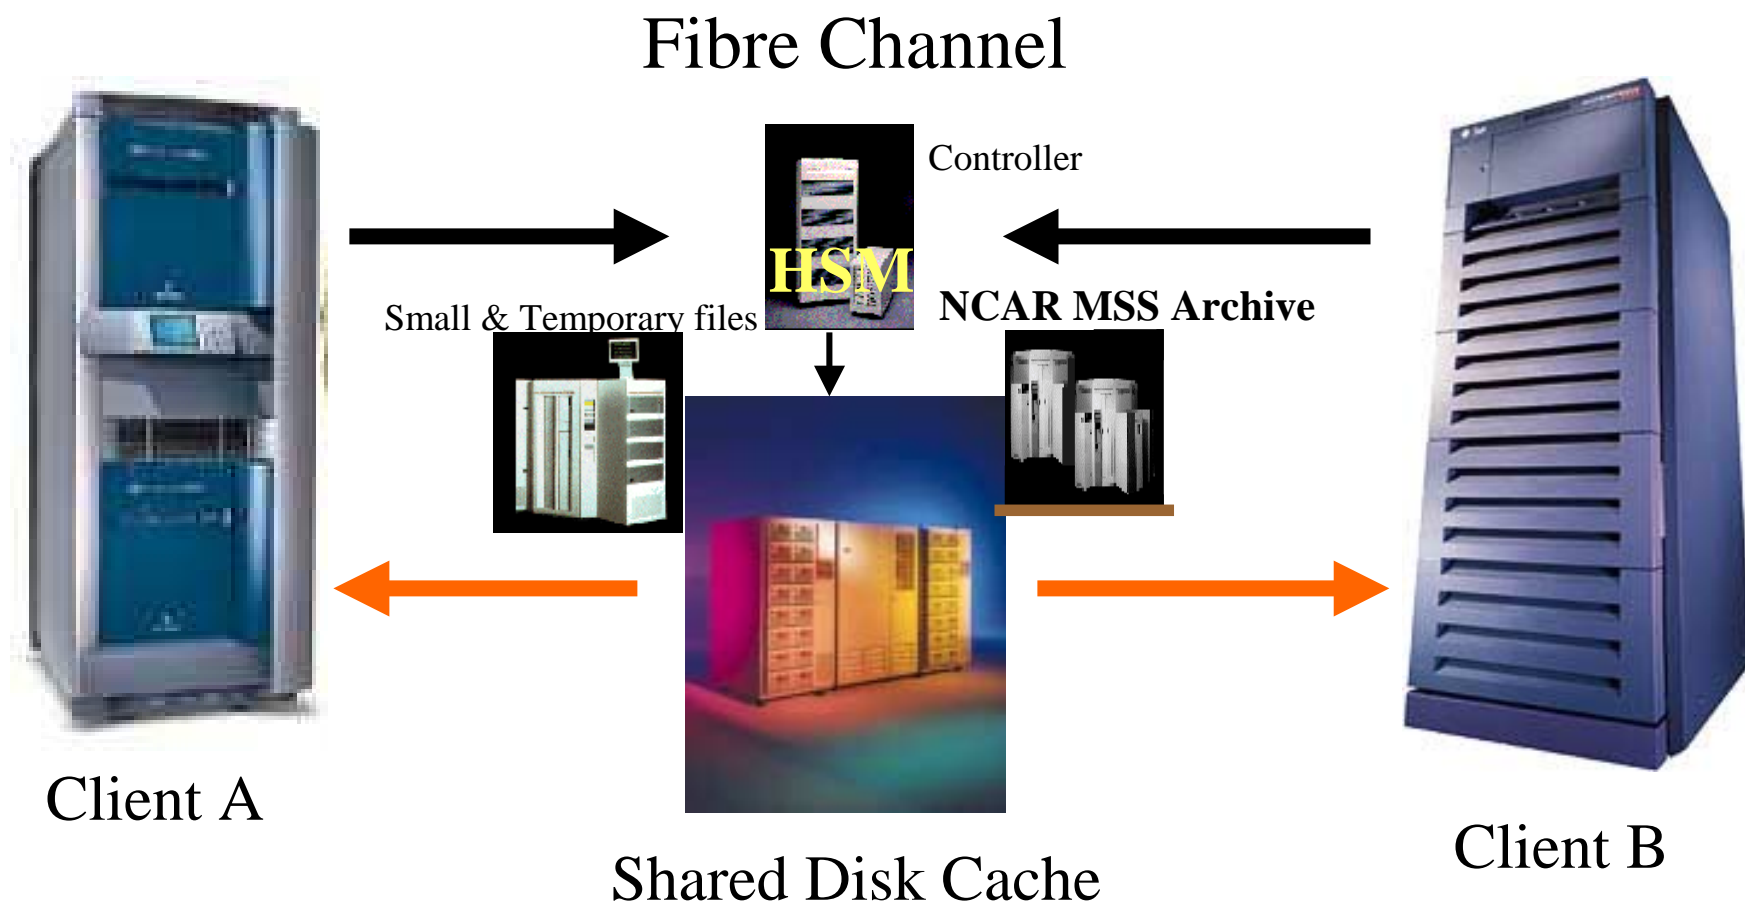


Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

16



Shared File System + HSM



Strategies II

■ Global Shared File Service

■ Better internal management

- Caching of small & temporary files
- Vendor manages software on heterogeneous servers (?)

■ Better access model

- Familiar POSIX interface (but no extended metadata)
- Record-level access
- Greater visibility (e.g., via NFS, SAMBA)
- GUI interfaces are already out there (e.g., Windows, mc)




Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

18



Strategies II: Catalog/Search Service



Strategies II

■ Catalog/Search Service

■ More and better extended metadata

- Restores what was lost to POSIX file service
- Metadata for data interpretation/extraction

■ Powerful search capabilities

- Domain-specific views
- Data subsetting

■ Better Access Model

- Similar to web search engine
- Consolidation of independent efforts across the organization
- Selective visibility of data to the general public
- Data delivery by appropriate (cost effective) means



Summary

- The prospect of overwhelming growth will require unprecedented changes to the NCAR Mass Storage System
- The rapidly changing computer/web landscape compounds the problem, but also offers solutions
- We are moving slowly towards storage solutions that can deliver meaningful information with less effort to more users with better “access models”



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

21



Summary



"...Things fall apart; the centre cannot hold..."

*"...And what rough beast, its hour come round at last,
Slouches towards Bethlehem to be born?"*

W.B. Yeats



Copyright © 2000 University Corporation for Atmospheric Research
THIC Meeting June 28, 2000

22

