Applications Drivers For Data "Parking" On The Information Superhighway

Clark E. Johnson, Jr.
Research Program on Communications Policy
Massachusetts Institute of Technology
P.O. Box 50116
Minneapolis, MN 55405
clark@farnsworth.mit.edu
Tel: 612-922-8541
Fax: 612-922-8820

Thomas Foeller
Managing Director
Configured Health Care, Inc.
RR1, Box 111
Balsam lake, WI 54180
71732.243@compuserv.com

Abstract

As the cost of data storage continues to decline (currently about one-millionth of its cost four decades ago) entirely new application areas become economically feasible. Many of these new areas involve the extraordinarily high data rates and universal connectivity soon to be provided by the National Information Infrastructure (NII).

The commonly held belief is that the main driver for the NII will be entertainment applications. We believe that entertainment applications as currently touted--multi-media, 500 video channels, video-on-demand, etc.--will play an important but far from dominant role in the development of the NII and its data storage components. The most pervasively effective drivers will be medical applications such as telemedicine and remote diagnosis, education and environmental monitoring. These applications have a significant funding base and offer a clearly perceived opportunity to improve the nation's standard of living.

The NII's wideband connectivity both nationwide and worldwide requires a broad spectrum of data storage devices with a wide-range of performance capabilities. These storage centers will be dispersed throughout the system. Magnetic recording devices will fill the vast majority of these new data storage requirements for at least the rest of this century.

The storage needs of various application areas and their respective market sizes will be explored. The comparative performance of various magnetic technologies and competitive alternative storage systems will be discussed.

OVERVIEW

Evolving local and wide-area networks are opening and supporting increasingly wider interoperation and data sharing. A number of these architectures also support real-time interoperable applications. Examples include on-line interactive games played over the Internet.

New commercial ventures are subjecting Internet facilities to increasing numbers of users with a wide variety of knowledge and skill levels. In addition, certain forthcoming Internet-based offerings (e.g. Delphi) will routinely use "agents" to traverse the Internet. These "agents" will be able to call up remote resources that can be aggregated to produce and deliver the desired result transparently as a "single" transaction.

Often times, valuable data artifacts resulting from or derived during these "aggregated" agent-driven operations are left at remote nodes for periods of time transcending the life of the transaction itself. These may create problems for both data owners and site operators where the data resides. Site operators will need to flush out such data from their storage devices from time to time, causing potential losses, without recourse, to the data owners.

To some degree, this same problem is already occurring at an increasing rate on smaller, local domains. Enterprise-wide networking that supports interoperable desktop computing often exhibits a potential soft underbelly of potential data loss. This is because there is no consistent nor convenient way to back-up widely dispersed data.

While there are currently a number of initiatives to develop and refine data-compacting interoperability models, there are none that we are aware of that are addressing back-up of widely dispersed data often typified by incompatible formats and inconsistent access mechanisms.

As much higher data rate networks such as the Global Information Infrastructure (GII) are developed that operate at speeds that are expected to approach a gigabit/second, the problems outlined above become intractable with current storage technology and system architecture. In order for such networks to become consistently useful and fulfill the promise of universal access to all, requires that dispersed data, located at nodes anywhere in the world be seamlessly available.
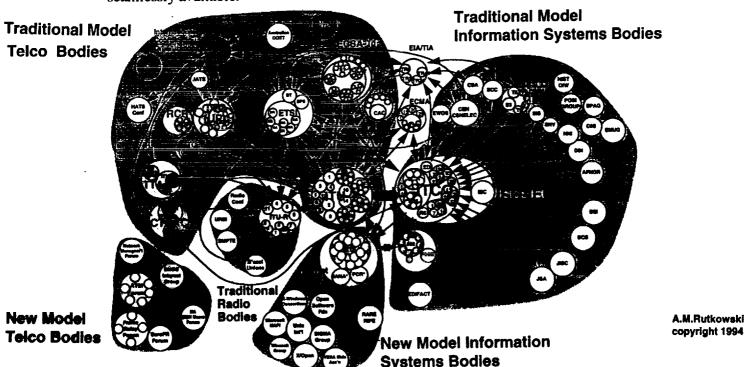

**Traditional Model
Telco Bodies**

**Traditional Model
Information Systems Bodies**

**New Model
Telco Bodies**

**Traditional
Radio
Bodies**

**New Model Information
Systems Bodies**

A.M.Rutkowski
copyright 1994

Figure 1. The Standards-Making Universe

2

This paper is a "first step" toward defining and formalizing a "universal" storage and retention model. We hope that this will encourage dialogue regarding the broad issues of the requirements of a truly user-friendly network, and to encourage the appropriate standards-setting groups to become actively involved. International standards setting is a convoluted process as Figure 1 shows.

SOME RELEVANT DRIVERS

Considerable effort is being made on a number of fronts to deal with the entire area of interoperability across dissimilar networks, user interfaces and underlying protocols. As these developments become commercialized, they will have a profound effect on the nature and structure of storage. Examples of the development efforts follow.

Most network management systems protocols and descriptive semantics, e.g. NMS (Novell), Openview (HP), Netmanager (SUN), Landesk (INTEL), Netview (IBM) do not currently interoperate across public network facilities such as the Internet, but they do operate across a number of commercial local and wide-area networks. Work is going on to provide upgrades that will soon enable certain user-specific backup and maintenance functions across heterogeneous networks.

Developments such as Microsoft "Windows '95" are expected to provide universal user interfaces including X-Windows along with the simultaneous support of multiple network transport mechanisms (e.g. TCP/IP, NetBIOS, netBEUI,...).

There are a number of applications program initiatives that will vie to become a "standard," and will enable application program modules written in one language to interoperate to some degree with programs written using one or more modules written in as many languages. These could, for example, share run-time resources in a single memory image; they could be segmented--running in separate memory images on different machines; or they might be distributed system components implemented by a number of programs on many machines. All of these implementations would appear seamless to the user; in fact, users may not know which were implemented and stored locally or at a remote node. As a consequence, slow or unreliable network transport mechanisms, inadequate long-term storage facilities located at network nodes, or missing software/data modules removed because of conflicting space requirements will increase and exacerbate real and perceived problems. Of course, increasing the number of interoperable components or modules expected by the user to be immediately available, will further increase the demand for high-performance storage.

BACKUP ISSUES

The above considerations lead to a number of data backup issues. As an example, consider the implementation of a flexible communication network backup facility that supports large block transfers of variable size and that functions without impacting the network's perceived performance by its on-line users. This is particularly important for digital video applications where even short time delays are intolerable.

In summary, because of this and myriad other issues, the overall storage architecture requires logical data space management so that it appears physically unbounded--regardless of the nature of the storage devices or their network topology. That is, object boundaries span storage hierarchies and media groupings. Consistent appearing data output formats are required so that the presentations are independent of origin data base, operational software and transport and are fully controllable by the user.

3

In addition, "agent" specifications and qualifications must be codified as there are compelling reasons to believe that these agents will be used to handle event log monitoring. billing (for example when the network is used to distribute intellectual property), reporting and internal data migration pathway control.

THE HIGH-SPEED NETWORK

Called the National Information Infrastructure in the United States, and the GII elsewhere, the vision is to have a "fat pipe" with Terabits/second bandwidth connecting major world centers by the next millennium. Currently, single-mode fiber-optic cables span the Atlantic and Pacific and operate at hundreds of megabits/second speeds. Terabytes of data per day will be traveling on this network and will require store-and-forward and destination node storage. Storage periods required will range from milliseconds to days.

Currently, Broadband Integrated Services Digital Network (BISDN), using the Synchronous Optical Network (SONET) physical layer protocol (SONET OC-3 operates at 155Mb/sec.) appears closest to commercialization, but requires three technological advances: in access, switching and storage.

In the area of storage, teleconferencing and multimedia applications such as MOSAIC can be accommodated on ISDN using commercialized bandwidth compression schemes. The storage requirements for applications such as medical image record transfer, where compression may not be acceptable to the profession, and for high-definition television are almost overwhelming.

For example, using the Qualcomm HDTV compression algorithm that provides 48:1 intra-frame compression, a full-length movie requires about 50 megabytes of storage. The average video store may have 10,000 titles in stock. To duplicate this on the server requires about 500 Terabytes of storage. These titles while not being required to be on one server, must be accessible under the same naming conventions, using consistent accessing and transfer methods, as if they were co-resident. Conventions and software required to meet these types of requirements are not yet available.

Missing or ill-defined technologies are not the only problems that must be faced. The economics of "video-on-demand" (VOD) is not comforting. The top one percent of Blockbuster customers spend $350/year. The average is about $50/year. Historically, citizens have spent a relatively fixed percentage of their income on entertainment--about 5% of their net after taxes; and this includes travel, restaurant meals and the like. It is questionable if, with currently available technology, there is a profitable business in downloading movies. The future availability of advanced, lower cost, faster storage means may well determine the fate of video on demand.

Medical image transmission and other telemedicine components are not yet inhibited by the same cost constraints as VOD. The pictures, which range from X-rays to magnetic-resonance images and sonograms, are sent without compression at relatively high cost because the medical profession (not to mention its regulators and the legal profession) are not yet willing to accept images reconstructed from image data that had been compressed using "lossy" compression algorithms--even though it can be shown that trained observers can rarely distinguish between the original images and those that have been compressed. Regardless of absolute cost, the value of using this technology to provide expert diagnosis to remote and rural areas cannot be overstated. Additionly, telemedicine using actual, high-resolution video

images provides the opportunity to perform almost real-time consultation and cooperative surgical procedures across the network.

America's schools are slowly becoming connected to each other and to world-wide data bases. While each institution's network usage, even if in multimedia format, will not put a significant demand on the network; the contemporaneous demand by many thousands of users will. A significant increase in response time when surfing the network for information using, for example Mosaic, will tend to discourage the students from using of the network's vast resources.

THE FUTURE OF STORAGE FOR THE HIGH-SPEED NETWORK

Certainly for the rest of this century, magnetic storage will dominate. The storage density continues to increase at the four decade-long historic rate of doubling every 2.5 years. In fact, this rate-of-change has recently increased. The per-bit cost of magnetic storage has gone down by a factor of over a million during the same time period. The retail price of disk storage is now roughly 40 cents/megabyte.

The always-increasing I/O problem (the mismatch between computational speed and memory access) is constantly being addressed by memory suppliers. The development of RAID (redundant array of inexpensive drives) disk technology and stripe-based tape storage provide performance and reliability improvements. Massive RAID systems now being developed are currently limited by the speed of silicon technology.

The usual hierarchy of storage will continue to prevail with the ultra-high speed cache requirements being fulfilled by solid-state memory. Massive caching is an economic issue because the cost of such memory is prohibitive for high data rate network applications. Thus, extensive effort is being devoted to refining and improving holographic storage technologies. Many researchers expect these to be commercialized by the end of the century and will combine the speed performance of RAM storage with the low-cost and capacity of magnetic-based systems. A particular attractiveness is the lack of moving parts and storage densities in the range of terabytes/cc.