STORIES

D-Lib Magazine
March 2000

Volume 6 Number 3

ISSN 1082-9873

Collection-Based Persistent Digital Archives - Part 1

Reagan Moore, Chaitan Baru, Arcot Rajasekar, Bertram Ludascher, Richard Marciano, Michael Wan, Wayne Schroeder, and Amarnath Gupta
<rmoore, baru, sekar, ludaesch, marciano, mwan, schroede, gupta@sdsc.edu>
San Diego Supercomputer Center

[This is the first of a two-part story. The second part will appear in the April 2000 issue of D-Lib Magazine.]

Abstract

The preservation of digital information for long periods of time is becoming feasible through the integration of archival storage technology from supercomputer centers, data grid technology from the computer science community, information models from the digital library community, and preservation models from the archivist's community. The supercomputer centers provide the technology needed to store the immense amounts of digital data that are being created, while the digital library community provides the mechanisms to define the context needed to interpret the data. The coordination of these technologies with preservation and management policies defines the infrastructure for a collection-based persistent archive [1]. This paper defines an approach for maintaining digital data for hundreds of years through development of an environment that supports migration of collections onto new software systems.

1. Introduction

Supercomputer centers, digital libraries, and archival storage communities have common persistent archival storage requirements. Each of these communities is building software infrastructure to organize and store large collections of data. An emerging common requirement is the ability to maintain data collections for long periods of time. The challenge is to maintain the ability to discover, access, and display digital objects that are stored within an archive, while the technology used to manage the archive evolves. We have implemented an approach based upon the storage of the digital objects that comprise the collection, augmented with the meta-data attributes needed to dynamically recreate the data collection. This approach builds upon the technology needed to support extensible database schema, which in turn enables the creation of data handling systems that

interconnect legacy storage systems.

The long-term storage and access of digital information is a major challenge for federal agencies. The rapid change of technology resulting in obsolescence of storage media, coupled with the very large volumes of data (terabytes to petabytes in size) appears to make the problem intractable. The concern is that when the data storage technology becomes obsolete, the time needed to migrate to new technology may exceed the lifetime of the hardware and software systems that are being used. This is exacerbated by the need to be able to retrieve information from the archived data. The organization of the data into collections must also be preserved in the face of rapidly changing database technology. Thus each collection must be migrated forward in time onto new data management systems, simultaneously with the migration of the individual data objects onto new media. The ultimate goal is to preserve not only the bits associated with the original data, but also the context that permits the data to be interpreted. In this paper we present a scalable architecture for managing media migration, and an information model for managing migration of the structure of the context. The information model includes a logical schema for organizing attributes, a physical characterization for how to load the attributes into the database, and a data dictionary for defining semantics.

We rely on the use of collections to define the context to associate with digital data. The context is defined through the creation of semi-structured representations for both the digital objects and the associated data collection. Each digital object is maintained as a tagged structure that includes either the original bytes of data or persistent links to the object, as well as attributes that have been defined as relevant for the data collection. The collection context is defined through use of both logical and physical representations for organizing the collection attributes. By using infrastructure independent representations, the original context for the archived data can be maintained. A collection-based persistent archive is therefore one in which the organization of the collection is archived simultaneously with the digital objects that comprise the collection [1].

A persistent collection requires the ability to dynamically recreate the collection on new technology. For a solution, we consider the integration of scalable archival storage technology from supercomputer centers, infrastructure independent information models from the digital library community, and preservation models from the archivist's community. An infrastructure that supports the continuous migration of both the digital objects and the data collections is needed. Scalable archival storage systems are used to ensure that sufficient resources are available for continual migration of digital objects to new media. The software systems that interpret the infrastructure independent representation for the collections are based upon generic digital library systems, and are migrated explicitly to new platforms. In this approach, the original representation of the digital objects and of the collections does not change. The maintenance of the persistent archive is then achieved through application of archivist policies that govern the rate of migration of the objects and the collection instantiation software.

The goal is to preserve digital information for at least 400 years. This paper examines the technical issues that must be addressed and presents a prototype implementation. The paper is organized into sections to provide a description of the persistence issues, and a generic description of the technology. (A description of the creation of a one million message persistent E-mail collection will be discussed in Part 2 in next month's issue of *D-Lib Magazine*.)

2. Persistence Issues

The preservation of the context to associate with digital objects is the dominant issue for collection-based persistent archives. The context is traditionally defined through specification of attributes that are associated with each digital object. The context is organized through relationships that exist between the attributes, and a description of the preferred organization of the attributes within user interfaces for accessing the data collection. We identify three levels of context that must be preserved:

- Digital object representation. Every digital object has attributes that define its structure, physical context, and provenance, and annotations that describe features of interest within the object. Since the set of attributes (such as annotations) will vary across all objects within a collection, a semi-structured representation is needed. Not all digital objects will have the same set of associated attributes.

- Data collection representation. The collection also has an implied organization, which is typically a subset of the attributes associated with the digital objects. A schema is used to support relational queries of the attributes or meta-data. It is possible to reorganize a collection into multiple tables to improve access by building new indexes, and in the more general case, by adding attributes. The schema used to organize the collection attributes can be different from the set of attributes associated with a digital object within the collection.

- Presentation representation. The user interface to the collection can present an organization of the collection attributes that is tuned to meet the needs of a particular community. Researchers may need access to all of the meta-data attributes, while students are interested in a subset. The structure used to define the user interface again can be different from the schema used for the collection organization. Each of these presentations represents a different view of the collection. Re-creation of the original view of a collection is a typical archival requirement.

Digital objects are used to encapsulate each data set. Collections are used to organize the context for the digital objects. Presentation interfaces are the structure through which collection interactions are defined. The challenge is to preserve all three levels of context for each collection.

2.1 Managing Context

Management of the collection context is made difficult by the rapid change of technology. Software systems used to manage collections are changing on three to five-year time scales. It is possible to make a copy of a database through a vendor specific dump or backup routine. The copy can then be written into an archive for long term storage. This approach fails when the database is retrieved from storage, as the database software may no longer exist. The archivist is then faced with migration of the data collection onto a new database system. Since this can happen for every data collection, the archivist will have to continually transform the entire archive. A better approach is needed.

An infrastructure independent representation is required for the collection that can be

maintained for the life of the collection. If possible, a common information model should be used to reference the attributes associated with the digital objects, the collection organization, and the presentation interface. An emerging standard for a uniform data exchange model is the eXtended Markup Language (XML) [2]. XML is the predominant instance of a semi-structured information model (i.e., labeled, ordered trees) and provides a representation for tagging data. Data could be relational data, object oriented data, schemas, procedures, etc. We define a collection as an XML view on the original tagged data using an information model. A particular example of an information model is the XML Document Type Definition (DTD) which provides a description for the allowed nesting structure of XML elements. Richer information models are emerging such as XSchema [3] (which provides data types, inheritance, and more powerful linking mechanisms) and XMI [4] (which provides models for multiple levels of data abstraction).

We shall reference the next generation information model as an Open Schema Definition (OSD). The OSD contains the collection schema and the presentation definition. For example, an XSL style sheet can be used for the presentation component of an OSD. XSL is the eXtensible Style sheet Language [5] and supports transformation of XML documents and formatting of the output for presentation. For the prototype we used XML DTDs and XSL style sheets together as the OSD.

It is possible to provide multiple presentational views of a collection. For our prototype we use multiple XSL style sheets for a specific collection to accommodate different user interfaces. The use of OSDs gives us freedom of choice for assembling a collection from tagged data objects, and for presenting the derived collection to multiple user communities.

Although XML DTDs were originally applied to documents only, they are now being applied to arbitrary digital objects, including the collections themselves. More generally, OSDs can be used to define the structure of digital objects, specify inheritance properties of digital objects, and define the collection organization and user interface structure.

While XML DTDs provide a tagged structure for organizing information, the semantic meaning of the tags is arbitrary, and depends upon the collection. A data dictionary is needed for each collection to define the semantics. A persistent collection therefore needs the following components of an OSD to completely define the collection context:

- Data dictionary for collection semantics,
- Digital object structure,
- Collection structure,
- User interface structure.

2.2 Managing Persistence

Persistence is achieved by providing the ability to dynamically reconstruct a data collection on new technology. While the software tools that do the reconstruction have to be ported to each new hardware platform or database, the collection can remain in its infrastructure independent format within an archive. The choice of the appropriate standard for the information model is vital for minimizing the support requirements for a collection-based persistent archive. The goal is to store the digital objects comprising the collection and the collection context in an archive a single time. This is possible if any

changes to the standard information model are added as a superset to the prior information model. The knowledge required to manipulate a prior version of the information model can then be encapsulated in the software system that is used to reconstruct the collection. With this caveat, the persistent collection never needs to be modified, and can be held as infrastructure independent bit-files in an archive.

The re-creation or instantiation of the data collection is done with a software program that uses the schema descriptions that define the digital object and collection structure to generate the collection. The goal is to build a generic program that works with any schema description. This will reduce the effort required to support dynamic reconstruction of a persistent data collection to the maintenance of a single software system.

Maintaining persistent digital objects requires the ability to migrate data to new media. The reasons for continuing to refresh the media on which the collection is maintained are:

- Avoid loss of data because of the finite lifetime and resulting degradation of the media.

- Minimize storage costs. New media typically store at least twice as much data as the prior version, usually at the same cost per cartridge. Thus migration to new media results in the need for half as many cartridges, decreased floor space, and decreased operating costs for managing the cartridges. Note that for this scenario, the media costs for a continued migration will remain bounded, and will be less than twice the original media cost. The dominant cost to support a continued migration onto new media is the operational support needed to handle the media.

- Maximize the ability to handle exponentially increasing data growth. Many data collections are doubling in size in time periods shorter than a year. This means the effort to read the entire collection for migration to new media will be less than the effort to store the new data that is being collected within that year. Migration to higher density media that has a faster read/write rate is the only way to guarantee the archived data is accessible. The governing metric for a collection is the total time required to re-read the entire collection. Unless the re-read time remains bounded, a persistent archive will become unmanageable.

To facilitate migration and access, supercomputer centers keep all data in tape robots. For currently available tape (cartridges holding 20 GB to 50 GB of data), a single tape robot is able to store 120 terabytes to 300 terabytes of uncompressed data. By year 2003, a single tape robot is expected to hold 6000 terabytes, using 1-terabyte capacity cartridges. The storage of petabytes (thousands of terabytes) of data is now feasible. The capacity of archives will not be a limiting factor.

Given that the collection context and the digital objects can be migrated to new media, the remaining system that must be migrated is the archival storage system itself. The software that controls the tape archive is composed of databases to store the storage location and name of each data set, logging systems to track the completion of transactions, and bitfile movers for accessing the storage peripherals. Of these components, the most critical resource is the database or nameserver directory that is used to manage the names and locations of the data sets. At the San Diego Supercomputer Center, the migration of the nameserver directory to a new system has been done twice, from the DataTree archival

storage system to the UniTree archival storage system, and from UniTree to the IBM High Performance Storage System [6]. Each migration required the read of the old directory, and the ingestion of each data set into the new system. Although the number of files increased from 4 million to 7 million between the two migrations, the time required for the migration decreased from 4 days to 1 day. This reflects advances in vendor supplied systems for managing the name space. Based on this experience, it is possible to migrate to new archival storage systems, without loss of data.

One advantage of archival storage systems is their ability to manage the data movement independently from the use of the data. Each time the archival storage system was upgraded, the new version of the archive was built with a driver that allowed tapes to be read from the old system. Thus migration of data between the archival storage systems could be combined with migration onto new media, minimizing the number of times a tape had to be read.

The creation of a persistent collection can be viewed as the design of a system that supports the independent migration of each internal hardware and software component to new technology. Management of the migration process then becomes one of the major tasks for the archivist.

2.3 Managing Scalability

A persistent archive can be expected to increase in size through either addition of new collections, or extensions to existing collections. Hence the architecture must be scalable, supporting growth in the total amount of archived data, the number of archived data sets, the number of digital objects, the number of collections, and the number of accesses per day. These requirements are similar to the demands that are placed on supercomputer center archival storage systems. We propose a scalable solution that uses supercomputer technology, based on the use of parallel applications running on parallel computers.

A scalable system is built by identifying both the capabilities that are best provided by each component, and the constraints that are implicit within each technology. Interfaces are then constructed between the components to match the data flow through the architecture to the available capabilities. Archival storage systems are used to manage the storage media and the migration to new media. Database management systems are used to manage the collections. Web servers are used to manage access to the system.

Archival storage systems excel at storing large amounts of data on tape, but at the cost of relatively slow access times. The time to retrieve a tape from within a tape silo, mount the tape into a tape drive, and ready the tape for reading is on the order of 15-20 seconds for current tape silos. The time required to spin the tape forward to the position of the desired file is on the order of 1-2 minutes. The total time can be doubled if the tape drive is already in use. Thus the access time to data on tape can be 2-4 minutes. To overcome this high latency, data is transferred in large blocks, such that the time it takes to transfer the data set over a communication channel is comparable to the access latency time. For current tape peripherals which read at rates from 10 MB/sec to 15 MB/sec, the average data set size in an archive should be on the order of 500 MB to 1 GB. Since digital objects can be of arbitrary size, containers are used to aggregate digital objects before storage into the archive.

The second constraint that must be managed for archives is the minimization of the number of data sets that are seen by the archive. Current archival storage nameservers are able to manage on the order of 10 - 40 million data sets. If each data set size is on the order of 500 MB, the archive can manage about 10 petabytes of data (10,000 TBs, or 10 million GBs). Archival storage systems provide a scalable solution only if containers are used to aggregate digital objects into large data sets. The total number of digital objects that can be managed is on the order of 40 billion, if one thousand digital objects are aggregated into each container.

Databases excel at supporting large numbers of records. Note that the Transaction Processing Council D benchmark [7] measures performance of relational databases on decision support queries for database sizes ranging from 1 gigabyte up to 3 terabytes and from 6 million to 18 billion rows. Each row can represent a separate digital object. With object relational database systems, a binary large object or BLOB can be associated with each row. The BLOBs can reside either internally within the database, or within an external file system. In the latter case, handles are used to point to the location of BLOB. The use of handles makes it feasible to aggregate digital objects within containers. Multiple types of container technology are available for aggregating digital objects. Aggregation can be done at the file level, using utilities such as the TAR program, at the database level through database tablespaces, or at an intermediate data handling level through use of software controlled caches. The database maintains the information needed to describe each object, as well as the location of the object within a container and the location of the container within the storage system. A data handling system is used to support database access to archival storage.

Queries are done across the attributes stored within each record. The time needed to respond to a query is optimized by constructing indexes across the database tables. This can reduce the time needed to do a query by a factor of a thousand, at the cost of the storage space for the index, and the time spent in assembling the index. Persistent collections may be maintained on disk to support interactive access, or they may be stored in the archive, and rebuilt on disk when a need arises. If the collection is reassembled from out of the archive, the dominant time needed for the process may be the time spent creating a new index. Since archival storage space is cheap, it may be preferable to keep both infrastructure independent and infrastructure dependent representations of a collection. The time needed to load a pre-indexed database snapshot is a small fraction of the time that it would take to reassemble and index a collection. The database snapshot, of course, assumes that the database software technology is still available for interpreting the database snapshot. For data collections that are frequently accessed, the database snapshot may be worth maintaining.

The presentation of information for frequently accessed collections requires Web servers to handle the user load. Servers function well for data sets that are stored on local disk. In order to access data that reside within an archive, a data handling system is needed to transfer data from the archive to the Web server. Otherwise the size of the accessible collection may be limited to the size of the Web server disk cache. Web servers are available that distribute their load across multiple CPUs of a parallel computer, with parallel servers managing over 10 million accesses per day.

Web servers provide a variety of user interfaces to support queries and information discovery. The preservation of the user interface requires a way to capture an

infrastructure independent representation for the query construction and information presentation. Web servers are available that retrieve information from databases for presentation. What is needed is the software that provides the ability to reconstruct the original view of the collection, based upon a description of the collection attributes. Such technology is demonstrated as part of the collection instantiation process in the SDSC persistent archive prototype.

2.4 Managing Heterogeneity of Data Resources

A persistent archive is inherently composed of heterogeneous resources. As technology evolves, both old and new versions of the software and hardware infrastructure will be present at the same time. An issue that must be managed is the ability to access data that is present on multiple storage systems, each with possibly different access protocols. A variant of this requirement is the ability to access data within an archive from a database that may expect data to reside on a local disk file system. Data handling systems provide the ability to interconnect archives with databases and with Web servers. Thus the more general form of the persistent archive architecture uses a data handling system to tie each component together. At the San Diego Supercomputer Center, a particular implementation of a data handling system has been developed, called the Storage Resource Broker (SRB) [8].

The SRB supports the protocol conversion needed for an application to access data within either a database, file system, or archive. The heterogeneous nature of the data storage systems is hidden by the uniform access API provided by the SRB. This makes it possible for any component of the architecture to be modified, whether archive, or database, or Web server. The SRB Server uses a different driver for each type of storage resource. The information for which driver to use for access to a particular data set is maintained in the associated Meta-data Catalog (MCAT) [9-10]. The MCAT system is a database containing information about each data set that is stored in the data storage systems. New versions of a storage system are accessed by a new driver written for the SRB. Thus the application is able to use a persistent interface, even while the storage technology changes over time.

3. Implementation Strategy

A collection-based persistent archive can be assembled using a scalable architecture. The scalable architecture relies upon parallel hardware and software technology that is commercially available. The persistent archive requires the integration of three separate components: archival storage, collection management, and access servers through the use of a data handling system. The result is a system that can be modified to build upon new technology on an incremental basis. For a persistent archive to work within this migration environment, the data context must be maintained in an information independent representation. The technology to instantiate the collection will have to be migrated forward in time, along with the data handling system. The collection can be kept as bit-files within the archive, while the supporting hardware and software systems evolve.

3.1 General Architecture

The implementation of a prototype persistent archive at SDSC is based upon use of commercially available software systems, augmented by application level software developed at the San Diego Supercomputer Center. The general architecture software

components are listed below, followed by the particular software system used for the prototype:

- Archival storage system - IBM High Performance Storage System (HPSS) [6]
- Data handling system - SDSC Storage Resource Broker (SRB) [8]
- Object relational database - Oracle version 7.3, IBM DB2 Universal Database
- Collection management software - SDSC Meta-data Catalog (MCAT) [9, 10]
- Collection instantiation software - SDSC scripts
- Collection ingestion software - SDSC scripts
- Semi-structured data model - eXtended Markup Language - Document Type Definition [2]
- Relational data model - ANSI SQL Data Definition Language [11]
- DTD manipulation software - UCSD XML Matching and Structuring language (XMAS) [12]
- Web server - Apache Web server
- Presentation system - Web Browser such as Internet Explorer version 5.

The hardware components are:

- Archival storage system - IBM SP 8-node, 32-processor parallel computer, 180 TB of tape storage, three Storage Technology tape robots, and 1.6 TB of RAID disk cache
- Data management system - Sun Enterprise 4-processor parallel computer
- Data ingestion platform - SGI workstation
- Network interconnect - Ethernet, FDDI, and HiPPI

Each of these systems is scalable, and can be implemented using parallel computing technology. The efficiency of the archival storage system is critically dependent upon the use of containers for aggregating data before storage. Three different mechanisms have been tried at SDSC:

- Unix utilities. The TAR utility can be used to aggregate files. For container sizes of 100 MB, the additional disk space required is minimal. The disadvantages are that the container must be read from the archive and unpacked before data sets are accessed.
- Database tablespace. At SDSC, a prototype version of the DB2 UDB [13] parallel

object-relational database has been used to support large data collections. The prototype database stores the digital objects internally within tablespaces. The tablespaces can be stored within the HPSS archival storage system, and retrieved to a disk cache on demand. This effectively increases the database storage capacity to the size of the archive, while simultaneously aggregating digital objects into containers before storage in the archive.

- Data handling software cache. The SDSC Storage Resource Broker supports containers. Digital objects that are written into an archive through the SRB are aggregated into a container on a disk cache. When the container is full, the SRB writes the container into the archive. When data is referenced, the container is retrieved from the archive and the data set is read directly out of the container by the SRB.

3.1.1 Archive

The core of the architecture is the archival storage system, as it ultimately determines the total capacity, data ingestion rate, and data migration support for the persistent archive. The High Performance Storage System (HPSS) is supported by a parallel computer, the IBM SP. HPSS at SDSC currently stores over 14 million files, with an aggregate size of 140 TB. Data movement rates have been achieved that exceed 1 TB of data storage per day. The system sustains 16,000 file operations per day. The HPSS system is accessed over high-speed networks through a High Performance Gateway Node (HPGN). The HPGN supports multiple types of network access, including a 100 MB/sec HiPPI network, 100 Mb/sec FDDI, and Ethernet. The HPGN is directly connected to the nodes of the SP on which the HPSS software system runs through the Trail Blazer 3 switch. The HPSS central control services run on one of the four-processor SP nodes, while the bitfile movers that read/write data off of disk and tape are distributed across seven of the SP nodes. By interconnecting the external networks through the HPGN onto the SP switch, all of the mover nodes can be used in parallel, sustaining high data throughput. By having disk and tape drives connected to each of the mover nodes, data can be migrated in parallel to tape. Measured data movement rates from the nodes to the HPGN are 90 MB/s for file sizes on the order of 10 MB.

The HPSS archive includes multiple backup systems for preserving the nameserver directory, including mirroring of the directory on disk, backup of snapshots of the directory onto tape, transaction logging of all changes to the directory, and reconciliation of the transaction logs with the directory snapshots on a daily basis. To handle disasters, copies of the critical data sets are maintained in a second HPSS archival storage system located within another city. A description of the backup systems is given in [14]. The attention paid to nameserver directory backup is of critical importance. If the nameserver directory is lost, it will not be possible to name the files stored in the archive.

The HPSS archive is scalable, through the addition of more nodes, disk, and tape drives. The system has recently been upgraded to a capacity of 360 GB of uncompressed data through the acquisition of tape drives that write 20 GBs of data per cartridge. The system supports data compression. For the scientific data sets stored at SDSC, the average compression ratio is a factor of 1.5, implying the total capacity of the system is 500 TB.

3.1.2 Data Handling System

The data handling system provides the ability to connect heterogeneous systems together. We provide a detailed description of the SDSC data handling system to illustrate the software infrastructure needed to provide location and protocol transparency. The data handling infrastructure developed at SDSC has two components: the SDSC Storage Resource Broker (SRB) [8] that provides federation and access to distributed and diverse storage resources in a heterogeneous computing environment, and the Meta-data Catalog (MCAT) [9] that holds systemic and application or domain-dependent meta-data about the resources and data sets (and users) that are being brokered by the SRB. The SRB-MCAT system provides the following capabilities:

- uniform APIs for access to heterogeneous file systems, databases, and archival storage,
- protocol-transparency and location-transparency when accessing distributed systems,
- uniform persistent name space abstraction [24] over the file systems that are being brokered,
- collection-based access to remote data sets, thus supporting information discovery based on domain and system-dependent meta-information stored along with (or extracted from) the stored files,
- facilities for replication, copying or moving files across heterogeneous systems, performing resource-level operations (proxy operations) on data before delivery to the client, and
- an integrated encryption and authentication system that can range from no security to fully encrypted and fully authenticated data transfer including security against man-in-the-middle security intrusions [15, 16].

The SDSC Storage Resource Broker (SRB) is middleware that provides distributed clients with uniform access to diverse storage resources in a heterogeneous computing environment. Storage systems handled by the current release of the SDSC SRB include the UNIX file system, archival storage systems such as UniTree, ADSM and HPSS, and database Large Objects managed by various DBMSs including DB2, Oracle, and Illustra. Currently, the system runs on supercomputers such as the CRAY C90, CRAY T3E and IBM SP, on workstations such as Sun, SGI, and Compaq platforms, and on Windows NT. The SRB API presents clients with a logical view of data sets stored in the SRB. Similar to the file name in the file system paradigm, each data set stored in SRB has a logical name, which may be used as a handle for data operation. Unlike the file system where the physical location of a file is implied in its path name through its mount point, the physical location of a data set in the SRB environment is logically mapped to the data sets. Therefore, data sets belonging to the same collection may physically reside in different storage systems. A client does not need to remember the physical mapping of a data set. It is stored as meta-data associated with the data set in the MCAT catalog. Data sets in the SRB are grouped into a logical (hierarchical) structure called *collections*. The collection provides an abstraction for:

- placing similar objects (possibly, physically distributed) under one collection (e.g., image collections of a museum) and
- placing all dissimilar objects that have a common connection under one abstraction (e.g., all the text paragraphs, images, figures, and tables of a document).

The SRB supports data replication in two ways. One can replicate an object during object creation or modification. To enable this, SRB and MCAT allow the creation of *logical storage resources* (LSR) which are a grouping of two or more resources. When an application creates or writes a data set to these logical resources, the operations are performed on each of the grouped resources. The result of using a LSR is that a copy of the data is created in each of the physical resources belonging to the logical resource. It is possible to specify that the write operation is successful if k of the n copies are created. The user can modify all the copies of the data by writing to the data set with a "write all." The SRB provides an off-line replication facility to replicate an existing data set. This operation can also be used for synchronization purposes. When accessing replicated objects, SRB will open the first available replica of the object as given by a list from MCAT. The SRB also provides authentication and encryption facilities [15, 16], access control list and ticket-based access [17], and auditing capabilities to give a feature-rich environment for sharing distributed data collections among users and groups of users.

The design of the SRB server is based on the traditional network connected client/server model but has the additional capability of federation. Once a connection from a client is established and authenticated, a SRB agent is created that brokers all the operations for that connection. A client application can have more than one connection to a SRB server and to as many servers as required. The federation of SRBs implies that a client connects to any SRB server while accessing a resource that is brokered by another server. An inter-SRB communication protocol supports the federation operation. The SRB communicates with MCAT to obtain meta-information about the data set, which it then uses for accessing the data set.

3.1.3 Collection Management

A characterization of a relational database requires a description of both the logical organization of attributes (the schema), and a description of the physical organization of attributes into tables. For the persistent archive prototype we used XML DTDs to describe the logical organization. The physical organization of relational databases was expressed using the Data Definition Language, DDL [11]. A combination of the schema and physical organization can be used to define how queries can be decomposed across the multiple tables that are used to hold the meta-data attributes. It is possible to generate arbitrary mappings between a DTD semi-structured representation, and a DDL relational representation of a collection. A preferred correspondence between the two representations must be defined if a relational database is used to assemble the collection. XML-based databases are becoming available that remove the need to describe the physical layout. Examples are Excelon [18] (an XML variant of ObjectStore) and Ariel [19] (an XML version of O2). By using an XML-based database, it is possible to avoid the need to map between semi-structured and relational organizations of the database attributes. This minimizes the amount of information needed to characterize a collection, and makes the re-creation of the database easier.

A detailed description of the SDSC MCAT system is provided to illustrate the complexity of the information management software needed to describe and manage collection level meta-data. The SDSC MCAT is a relational database catalog that provides a repository of meta information about digital objects. Digital object attributes are separated into two classes of information within the MCAT:

- System-level meta-data that provides operational information. These include information about resources (e.g., archival systems, database systems, etc., and their capabilities, protocols, etc.) and data objects (e.g., their formats or types, replication information, location, collection information, etc.).

- Application-dependent meta-data that provides information specific to particular data sets and their collections (e.g., Dublin Core [20, 21] values for text objects).

Both of these types of meta-data are extensible, i.e., one can add and/or remove attributes. Internally, MCAT keeps schema-level meta-data about all of the attributes that are defined. The schema-level attributes are used to define the context for a collection and enable the instantiation of the collection on new technology. The attributes include definition of:

- Logical Structure: When a set of meta-data is registered with MCAT, one needs to identify a logical structure in which the rest of the meta-data will be organized. The logical structure should not be confused with database schema and are more general than that. For example, we have implemented the Dublin Core database schema [20] to organize attributes about digitized text. The attributes defined in the logical structure that is associated with the Dublin Core schema contains information about the subject, constraints, and presentation formats that are needed to display the schema along with information about its use and ownership.

- Attribute Clusters: An attribute cluster is a set of attribute names that are logically interconnected and that have a one-to-one mapping among them. One can view them as a (single or a set of) normalized table(s) in a database context. For example, in the Dublin Core, publisher, name, address, and contact information form a cluster. Contributor name and contributor type form a second cluster; title and its type form yet another cluster, and so on. Similarly in our system-level MCAT core meta-data, we have one cluster for each data replica containing the type, location, and size of the data objects. This aids the implementation of relational joins across the meta-data tables, since each replica has only one value for these properties and these properties provide the physical characteristics of the object. For each cluster, MCAT keeps information about any constraints and comments that can be searched when using the attribute, along with information about use-privileges and grant-of-use-privileges for the cluster. For each attribute, MCAT keeps more than 20 different types of information including its physical, logical and input and output characteristics [9].

- Token Attributes: Token attributes have a specific function (compared to other attributes); they capture some simple semantic information about the domain of discourse. One can also use the token attribute to capture semantic translation between discipline domains (e.g., common names vs. scientific names) and also capture hierarchical and equivalence relationships in the domain of discourse. Given

the development of semantic standards within a discipline, one can use the token attribute as a bridge between two schemas and provide semantic interoperability.

- Linkages: Linkages provide a means for inter-operating within and between schema. One can define four types of linkages:

  1. attribute-to-attribute,
  2. cluster-to-attribute,
  3. cluster-to-cluster, and
  4. cluster-to-token.

  Each of the linkages can be from one-to-many, many-to-one, or many-to-many. The linkage information is used to generate joins dynamically based on the user's chosen set of attributes. The join algorithm uses Steiner Tree generation of SQL commands from a directed acyclic graph; the DAG is a mapping of clusters and the linkages between them. The linkage information is used to perform federated query operations across schemas. The DAG is also used to figure out the notion of an allowed query by disallowing queries that span disjointed graphs.

MCAT provides APIs for creating, modifying and deleting the above structures. MCAT provides an interface protocol for applications such as Web servers. The protocol uses a data structure for the information interchange which is called MAPS -- Meta-data Attribute Presentation Structure. The data structure, which also has a wire-format for communication and a data format for computation, provides an extensible model for communicating meta-data information. A mapping is being developed to translate from the MAPS structure to the Z39.50 format [22]. Internal to MCAT, the schema for storing meta-data (may possibly) differ from MAPS, and hence mappings between the internal format and MAPS are needed for every type of implementation of the MCAT. Note that it is possible to store the meta-data in databases, flat files, or LDAP directories [23]. MAPS provides a uniform structure for communicating between MCAT servers and user applications.

The MAPS structure defines a query format, an update format and an answer format. The MAPS query format is used by MCAT in generating joins across attributes based on the schema, cluster and linkages discussed above. Depending upon the internal catalog type (e.g., DB2 database, Oracle database, or LDAP) a lower-level target query is generated. Moreover, if the query spans several database resources, a distributed query plan is generated.

The MCAT system supports the publication of schemata associated with data collections, schema extension through the addition or deletion of new attributes, and the dynamic generation of the SQL that corresponds to joins across combinations of attributes. GUIs have been created that allow a user to specify a query by selecting the desired attributes. The MCAT system then dynamically constructs the SQL needed to process the query. By adding routines to access the schema-level meta-data from an archive, it is possible to build a collection-based persistent archive. As technology evolves and the software infrastructure is replaced, the MCAT system can support the migration of the collection to the new technology. Effectively, the collection is completely represented by the set of digital objects stored within the archive, the schema that contains the digital object meta-data, and the schema-level meta-data that allows the collection to be instantiated

from scratch.

To Be Continued:

The first part of this article has concentrated on a description of the persistence issues, and a generic description of the scalable technology for managing media and context migration. The second part of the article will describe the creation of a one million message persistent E-mail collection. It will discuss the four major components of a persistent archive system: support for ingestion, archival storage, information discovery, and presentation of the collection. The technology to support each of these processes is still rapidly evolving, and opportunities for further research are identified.

References

[1] Rajasekar, A., Marciano, R., Moore, R., "Collection-Based Persistent Archives," Proceedings of the 16th IEEE Symposium on Mass Storage Systems, March 1999.

[2] Extensible Markup Language, <http://www.w3.org/XML>

[3] XSchema - representation of XML DTDs as XML documents, <http://www.simonstl.com/xschema/>

[4] XMI - XML Metadata Interchange, <http://www.omg.org/cgi-bin/doc?ad/99-10-02>

[5] XSL - eXtensible Stylesheet Language, W3C working draft, March 2000, <http://www.w3.org/TR/xsl/>

[6] The High Performance Storage System (HPSS), <http://www.sdsc.edu/hpss/>.

[7] Transaction Processing Council, <http://www.tpc.org/results/tpc_d.results.page.html>

[8] Baru C., Moore, R., Rajasekar, A., and Wan, M., "The SDSC Storage Resource Broker," *Proceedings of CASCON'98 Conference*, Nov. 30-Dec. 3, 1998, Toronto, Canada.

[9] Baru C., Frost, R., Marciano, R., Moore, R., Rajasekar, A., and Wan, M., "Meta-data to support information based computing environments," *Proceedings of the IEEE Conference on Meta-data*, Silver Spring, MD, Sept. 1997.

[10] MCAT - A Meta Information Catalog (V1.1), Technical report: <http://www.npaci.edu/DICE/SRB/mcat.html>

[11] Data Definition Language standardized syntax, ANSI X3.135-1992 (R1998)

[12] Baru C., Gupta, A., Ludascher, B., Marciano, R., Papakonstantinou, Y., Velikhov, P., and Chu, V., "XML-Based Information Mediation with MIX", Proceedings SIGMOD, Philadelphia, 1999.

[13] The DB2/HPSS Integration project, <http://www.sdsc.edu/MDAS>.

[14] Moore, R., Lopez, J., Lofton, C., Schroeder, W., Kremenek, G., Gleicher, M., "Configuring and Tuning Archival Storage Systems," Proceedings of the 16th IEEE Symposium on Mass Storage Systems, March 1999.

[15] Schroeder W., "The SDSC Encryption / Authentication (SEA) System," Distributed Object Computation Testbed (DOCT) project white paper, <http://www.sdsc.edu/~schroede/sea.html>.

[16] Schroeder W., "The SDSC Encryption and Authentication (SEA) System," to be published in Special Issue of Concurrency: Practice and Experience-Aspects of Seamless Computing, John Wiley & Sons Ltd.

[17] Baru C., and Rajasekar, A., "A Hierarchical Access Control Scheme for Digital Libraries," *Proceedings of the 3rd ACM Conference on Digital Libraries*, Pittsburgh, PA, June 23-25, 1998.

[18] Excelon XML database, <http://www.odi.com/excelon/main.htm>

[19] Eric N. Hanson, "The Design and Implementation of the Ariel Active Database Rule System", IEEE Transactions on Knowledge and Data Engineering , Vol. 8, No. 1, February 1996

[20] The Dublin Core, <http://purl.oclc.org/dc/>.

[21] The Warwick Framework, Carl Lagoze, *D-Lib Magazine*, July/August 1996, <http://www.dlib.org/dlib/july96/lagoze/07lagoze.html>

[22] Tomer, C., "Information Technology Standards for Libraries," *Journal of the American Society for Information Science*. 43: 566-570, 1992.

[23] Light-Weight Directory Access Protocol (LDAP) implementation by Netscape, <http://www.umich.edu/~dirsvcs/ldap/>

[24] Persistent namespace abstraction, the Handle System <http://www.handle.net>.

DOI: 10.1045/march2000-moore-pt1

SERIES_____ SUBSERIES_____
BOX_____ FOLDER_____

FermiNews

Fermi National Accelerator Laboratory

Volume 19 | Friday, August 16, 1996 | Number 16

INSIDE

Detour Ahead

East Entrance to Close Temporarily

Fermilab will close the east entrance to the Laboratory from 6 a.m. Monday, August 19 through 8 p.m. Friday, August 23, to permit maintenance on the railroad grade crossing by the Elgin, Joliet and Eastern Railroad. During the temporary closing, employees and others who normally enter via Batavia Road on the east side of the Laboratory must enter instead via Pine Street on the west. The entrance will also be closed to bicycles and pedestrians.

Detour signs will reroute traffic via Route 59 south, Butterfield Road west, and Kirk Road north to Pine Street. Questions? Call the Office of Public Affairs at 630-840-3351.

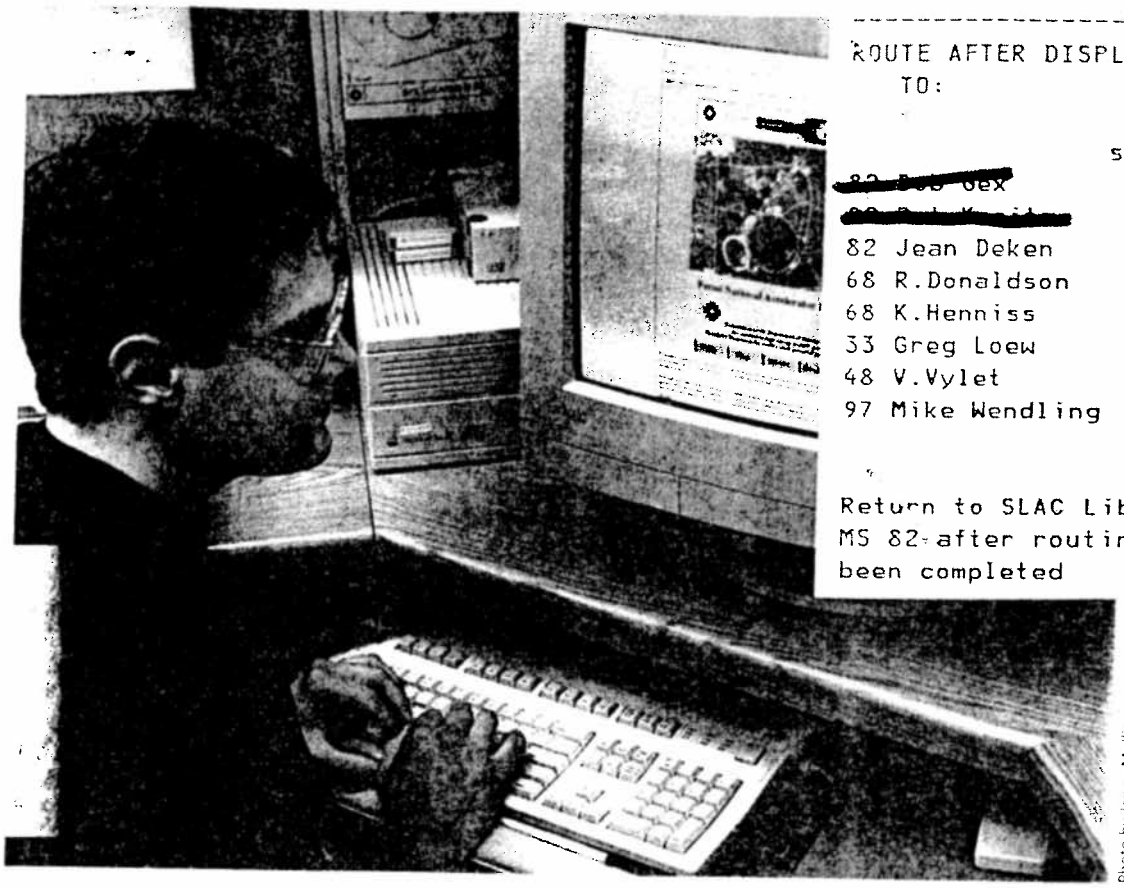

Photo by Jenny Mullins

Glenn Blanford, a graduate student from the University of California-Irvine, at work on the new Fermilab home page that will debut at the end of August.

High-Energy Physics

Birthplace of the Web

by Eric Berger, Office of Public Affairs

In mid-1991, when it spun its first thread from CERN to Fermilab, hardly anyone had ever heard of the World Wide Web. Now, only four years later, the Web is woven through the fabric of our culture, connecting people all over the planet in a new medium of communication. But of the millions who type "http://www...." each day, how many realize that—like the universe itself—it all began with high-energy physics?

What started out as tool for far-flung scientific colleagues to share each others' data may ultimately count among high-energy physics' most significant contributions to modern technology.

The Web was born in 1990, when Tim Berners-Lee, a computer scientist at CERN, the European Laboratory for Particle Physics, programmed the first types of computer codes, called protocols, that allowed a computer any-

continued on page 8

The Web at Fermilab

continued from page 7

Tim Berners-Lee, who wrote the software protocols to set up the World Wide Web in 1990 and 1991. He now directs the W3 Consortium, an open forum of companies and organizations with the mission to realize the full potential of the Web.

where to contact any other computer and communicate unfettered by log-in accounts and database incompatibilities.

"He's someone who thinks very clearly about the fundamentals, but besides this, he happened to come upon his idea at the right time," said Ruth Pordes, Computing Division Online Systems Department head. "His idea was to make distributed information over a network easily accessible to the whole world, and to work out the protocol of how to do that."

In its infancy, an internal laboratory version of the Web allowed high-energy physicists at CERN to share data and papers within their own network, but it remained difficult to connect to outside servers—from CERN to Fermilab for example—to share documents with other scientists.

Pordes, who had met Berners-Lee at international meetings, invited him to visit Fermilab for a few weeks in January, 1991. The visit coincided with a conference on hypertext, familiar to most Internet users as the "link" function of Web browsers. The conference gave Berners-Lee a stimulus for developing his ideas.

Berners-Lee returned to Fermilab in 1992. During his visit he made the first one-click link between the CERN server and Fermilab's central computers.

"We installed it (the CERN Web server) on FNALV and we linked it up to the documents used for data-taking and later analyses," Pordes said.

Fermilab's Jonathan Streets, leader of the Experiments Online Support Group, helped Berners-Lee set up the system.

"He and I wrote the server on FNALV that served the documents," Streets said. "Now anybody could come in and get them. That was the first time anybody could use the same interface to read documents pertaining to both data-taking and analyses—because we had very different databases."

Before Berners-Lee wrote the hypertext transfer protocols each computer server generally required a separate log-in account and had a different documentation system, making smooth communication between systems difficult. Berners-Lee's software (the Web), broke down the barriers dividing the hardware (the Internet).

The following summer, Berners-Lee gave the first United States presentation on the World Wide Web's merits.

"Ruth basically has to take credit for spotting the Web," said Vicky White, Computing Division deputy head. "We both knew Tim Berners-Lee's work, but she showed the first interest in it."

That Was Then...

Berners-Lee used these two transparencies in a 1992 Fermilab talk highlighting the World Wide Web.

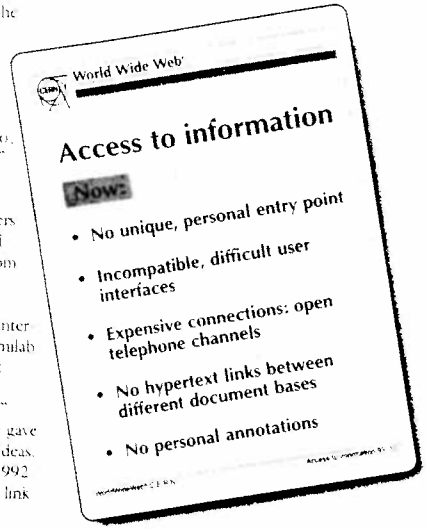

Fermilab joined the World Wide Web early. The Sloan Digital Sky Survey, an ambitious sky-mapping project in which Fermilab collaborates, became one of the first projects to adopt the Web in 1992.

"Sloan is a project that has a far-flung collaboration, and we felt having information that could be easily retrieved would be beneficial," Pordes said.

Beyond High-Energy Physics

Berners-Lee and others soon discovered the Web applied to areas well beyond high-energy physics.

"After that, he started working on protocols for the World Wide Web in order to have communication not only among physicists, but also to be able to find information from all the various other places available," Pordes said.

The next step required making the technology more accessible. Current Internet users surf the web using browsers such as Netscape or Web Explorer that offer windows into the Web. But in 1992 none of these mouse-friendly programs existed. A program called Mosaic paved the way for later browsers.

"Around 1992 NCSA (National Center for Supercomputing Applications at the University of Illinois at Urbana/Champaign) started work on Mosaic," Pordes said. "In 1993 I invited Tim to Fermilab. I contacted NCSA and suggested we go down and meet them. It was clear they were really going to market Mosaic. That was the first time we realized that he was really going to have a big success."

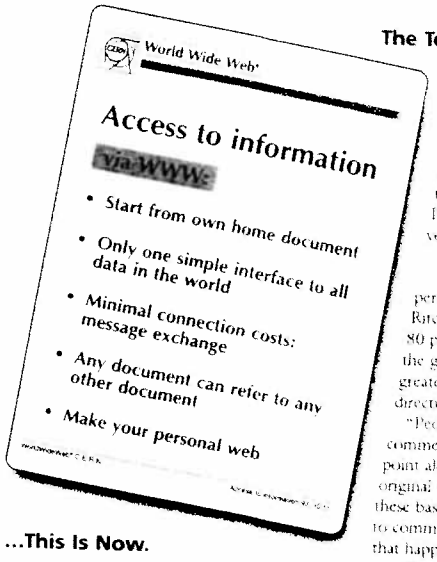

...This Is Now.

Mosaic brought many people outside high-energy physics into the Web.

"It wasn't until NCSA made Mosaic that the average person could see why this was really a good thing," said David Ritchie, who chaired the working group that designed the first official Fermilab home page.

Fermilab's role in the development of the Web came in bits and pieces.

"Fermilab filled one of several support roles," Pordes said. "There were people here who felt it was worthwhile at the beginning and were willing to put a little effort into it."

Pordes believes one of the reasons for the Web's phenomenal growth derives from its transatlantic origins, developing simultaneously in Europe and the United States.

"I think it's actually a big success partly because the Europeans felt this was their idea and they could adopt it and support it without thinking this was just something from America coming in and taking over," she said.

Another boost came when Stanford Linear Accelerator Center made SPIRES, a server that makes available journal article preprints, accessible to the Web, which attracted many scientists.

The Today and Tomorrow of The Web

What do long time users of the Web think of today's glossy incarnation?

"People complain about it a lot; they complain that all the information that's on it is junk and it's too slow," Pordes says. "I just say, 'Well, would you rather have it or not?'"

Clearly, most would rather have it.

Recent studies have found that eight percent of businesses have Web sites. Ritchie believes that figure will grow to 80 percent or more, but he sees a wane in the gold rush mentality and expects the greatest Web expansion to come from directions other than commercialism.

"People think the prime motivation is commercialism, but that's only a stopping point along the way," said Ritchie, the Lab's original webmaster. "It's really motivated by these basic kinds of desires by human beings to communicate to other human beings things that happen in their lives."

Growth on the individual level will come from new technologies that make access and manipulation of Web mechanisms easier. Ritchie feels that if such technology arises, the use of the Internet by individuals will skyrocket. Almost everyone who worked on the Web during its early years expresses surprise at its exponential growth.

"I think it's been absolutely amazing how quickly it's penetrated the average person's life," said White, of the Computing Division. "You never would have dreamt it three years ago."

Another interesting question remains the fascination new and continuing users have with the Web and the motivations that drive them to put many types of information online.

"There's a whole spectrum of motivations for people," White said. "It's probably the same reason people want to show you the photos of their grandchildren or their pictures of their vacation and to tell you about their work. It's a fundamental human need to share information. This is a tool that makes it easy to do that." ■

"It's a fundamental human need to share information. This is a tool that makes it easy to do that."

—Vicky White
Computing Division

Photo by Reg Jones

Computer scientist Jonathan Streets picks up Amy, 5, and Joseph, 2, from Fermilab Day Care. In 1992, Streets and Berners-Lee wrote the software that linked Fermilab to the World Wide Web.

Received: from SLACVM by SLACVM.SLAC.STANFORD.EDU (Mailer R2.08 R208004) with BSMTP id 0535; Mon, 09 Sep 96 18:26:00 PST
Received: from serv05.slac.stanford.edu by SLACVM.SLAC.STANFORD.EDU (IBM VM SMTP V2R1) with TCP; Mon, 09 Sep 96 18:25:59 PST
Received: from DIRECTORY-DAEMON by SERV05.SLAC.STANFORD.EDU (PMDF V5.0-6 #10979) id <01I9A98B4DQO00031S@SERV05.SLAC.STANFORD.EDU> for jmdeken@SLACVM.SLAC.Stanford.EDU; Mon, 09 Sep 1996 18:27:15 -0800 (PST)
Received: from [134.79.129.179] ("port 1474"@ROSELLMAC.SLAC.Stanford.EDU) by SERV05.SLAC.STANFORD.EDU (PMDF V5.0-6 #10979) id <01I9A988YTAC000301@SERV05.SLAC.STANFORD.EDU> for jmdeken@SLAC.Stanford.EDU; Mon, 09 Sep 1996 18:27:15 -0800 (PST)
Date: Mon, 09 Sep 1996 18:31:06 -0800
From: xanadu@SLAC.Stanford.EDU (P.A. Moore)
Subject: Re: First Web Server in US
X-Sender: xanadu@popserv.slac.stanford.edu
To: jmdeken@SLAC.Stanford.EDU
Message-id: <v0211010aae5a83548235@[134.79.129.179]>
X-Envelope-to: jmdeken@SLACVM.SLAC.Stanford.EDU
MIME-version: 1.0
Content-type: text/plain; charset="us-ascii"
Content-transfer-encoding: 7BIT

Whoops - a date mistake...see below.
PA

>From: cottrell@SLAC.Stanford.EDU
>Date: Mon, 09 Sep 1996 18:11:48 -0700
>Subject: Re: First Web Server in US
>X-Sender: cottrell@popserv.slac.stanford.edu
>To: brichter@SLAC.Stanford.EDU
>Cc: henniss@SLAC.Stanford.EDU, bebo@SLAC.Stanford.EDU,
> winters@SLAC.Stanford.EDU, xanadu@SLAC.Stanford.EDU,
> pfkeb@SLAC.Stanford.EDU, tony_johnson@SLAC.Stanford.EDU,
> leith@SLAC.Stanford.EDU, cottrell@SLAC.Stanford.EDU
>MIME-version: 1.0
>
>>Date: Mon, 09 Sep 1996 17:13 -0800 (PST)
>>From: "Joan M. Winters" <WINTERS@SLAC.Stanford.EDU>
>>Subject: Re: First Web Server in US
>>To: cottrell@SLAC.Stanford.EDU
>>
>>> The SLAC Web server was running on SLACVM (put up by Terry Hung) in November
>>> 1992.
>>
>>Don't you mean 1991? Joan
>
>Joan points out that I had a critical date wrong. Also the following is
>from Joan and of interest:
>
>
>Just in case you haven't bumped into it lately, here's Tony's
>Beam Line timeline, which is a tad different from his article:
>
> http://www-sld.slac.stanford.edu/sldwww/beamline/history.html
>
>When Paul Kunz brought the Web to SLAC, did he put up a server as
>well as a browser on the NeXT? Tony's January, 1991, date for SLAC
>is quite a lot earlier that what I heard from Terry Hung.
> Joan
>

>Les Cottrell
>Mail Stop 97, Stanford Linear Accelerator Center, POB 4349, Stanford CA 94309
>Phone: (415)926-2523, FAX: (415)926-3329
>WWW: http://www.slac.stanford.edu/~cottrell/
>

P.A. Moore, EdD
Assistant to the Director
Stanford Linear Accelerator Center
Box 4349 MS 80
Stanford CA 94309
phone 415-926-2605
fax 415-926-2525

Received: from SLACVM by SLACVM.SLAC.STANFORD.EDU (Mailer R2.08 R208004) with BSMTP id 0504; Mon, 09 Sep 96 18:10:54 PST
Received: from serv05.slac.stanford.edu by SLACVM.SLAC.STANFORD.EDU (IBM VM SMTP V2R1) with TCP; Mon, 09 Sep 96 18:10:54 PST
Received: from DIRECTORY-DAEMON by SERV05.SLAC.STANFORD.EDU (PMDF V5.0-6 #10979) id <01I9A8PL9ZG000042P@SERV05.SLAC.STANFORD.EDU> for jmdeken@SLACVM.SLAC.Stanford.EDU; Mon, 09 Sep 1996 18:12:10 -0800 (PST)
Received: from [134.79.129.179] ("port 1537"@ROSELLMAC.SLAC.Stanford.EDU) by SERV05.SLAC.STANFORD.EDU (PMDF V5.0-6 #10979) id <01I9A8PJ4WW4000301@SERV05.SLAC.STANFORD.EDU> for jmdeken@SLAC.Stanford.EDU; Mon, 09 Sep 1996 18:12:09 -0800 (PST)
Date: Mon, 09 Sep 1996 18:16:01 -0800
From: xanadu@SLAC.Stanford.EDU (P.A. Moore)
Subject: First Web Server in US
X-Sender: xanadu@popserv.slac.stanford.edu
To: jmdeken@SLAC.Stanford.EDU
Message-id: <v02110108ae5a7fbcaa2d@[134.79.129.179]>
X-Envelope-to: jmdeken@SLACVM.SLAC.Stanford.EDU
MIME-version: 1.0
Content-type: text/plain; charset="us-ascii"
Content-transfer-encoding: 7BIT

Jean- lots of us saw the article, including Burt, so I asked Les COttrell to do some research. Here is his response.
PA Moore

From: cottrell@SLAC.Stanford.EDU
>Date: Mon, 09 Sep 1996 17:16:08 -0700
>Subject: First Web Server in US
>X-Sender: cottrell@popserv.slac.stanford.edu
>To: brichter@SLAC.Stanford.EDU
>Cc: hennis@SLAC.Stanford.EDU, bebo@SLAC.Stanford.EDU,
> winters@SLAC.Stanford.EDU, xanadu@SLAC.Stanford.EDU,
> pfkeb@SLAC.Stanford.EDU, tony_johnson@SLAC.Stanford.EDU,
> leith@SLAC.Stanford.EDU, cottrell@SLAC.Stanford.EDU
>MIME-version: 1.0
>
>Thanks for the copy of the FermiNews for August 16, 1996. Looking at
>the headline and quickly browsing the article it leaves one with the
>feeling that FNAL was the pioneer site in the US. On more careful reading,
>they never make such a statement, the closest they get is to say
>
>Title:
>
>"High Energy Physics
>
>Birthplace of the Web"
>
>
>...in 1992 ...
>"He [Time Berners-Lee] and I [Jonathan Streets of FNAL] wrote the server
>on FNALV that served the documents. Now anybody could come in and get them.
>That was the first time anybody could use the same interface
>to read documents pertaining to both data-taking and analyses"
>
>I have talked to several people about this including Terry Hung, Kathryn
>Henniss, Bebo White (via Kathryn), and Joan Winters.
>

>The SLAC Web server was running on SLACVM (put up by Terry Hung) in November
>1992. Tim Berners-Lee anounced the Web server at info.cern.ch to the
>www-internet mailing list on Nov 12. In that announcement he said it "has been
>running for some time". The CERN Newsletter announced the Web to the world
>in Dec 1991.
>
>The official Web history at http://www.w3.org/pub/WWW/History.html
>says that SLAC was the first US server in March 1991. It also says for June
>1992 that FNAL joined with a Web server. Paul Kunz who might be able to
>spread more light is away at the moment, as is Tony Johnson.
>
>Bottom line, is that it is generally accepted that SLAC was the first Web
>server
>in the US, and FNAL do not deny this, though the article might lead one to
>think oterwise.
>
>Les Cottrell
>Mail Stop 97, Stanford Linear Accelerator Center, POB 4349, Stanford CA 94309
>Phone: (415)926-2523, FAX: (415)926-3329
>WWW: http://www.slac.stanford.edu/~cottrell/
>

P.A. Moore, EdD
Assistant to the Director
Stanford Linear Accelerator Center
Box 4349 MS 80
Stanford CA 94309
phone 415-926-2605
fax 415-926-2525

National Archives and Records Administration
Records Management

January 12, 2001

MEMORANDUM TO CHIEF INFORMATION OFFICERS:
Snapshot of Agency Public Web Sites

The National Archives and Records Administration (NARA) will be preserving a one-time snapshot of agency public web sites as they exist on or before January 20, 2001, as an archival record in the National Archives of the United States. The web site snapshot initiative outlined in this memorandum will ensure that we are able to document at least in part agency use of the Internet at the end of the Clinton Administration. To accomplish this, we ask that agencies take the following actions:

- Take a snapshot of your agency's public web site(s) following the instructions in the attachment to this memorandum.
- Within 60 days, send the snapshot and related documentation, also described in the attachment, to NARA.

I am sending this message to agency records officers also. I urge you to work with your records officer to ensure that your agency's public web sites are included in the snapshot.

This immediate action does not address the larger and long-term records management issues relating to agency web sites that you and NARA face. Agencies are using the web more and doing more of their business on the web than ever before. Web sites include different types of records that are used in different ways, and document different agency activities of varying significance. Some web content parallels traditional paper records. In fact some agencies use the web to disseminate policy information, but publicize the fact that the authentic records of policy are maintained on paper. But some web sites involve records that are important evidence; some are without precedent. Web sites cannot be treated with a "one size fits all" solution and Federal agencies need comprehensive guidance on management of these records.

We will be working with agencies on the first phase of this guidance in the next few weeks. We will meet with records officers shortly to go over what they need to see in the guidance. Several agencies have submitted records schedules for their web sites and have received dispositions for them. NARA has also provided training and briefings on records management issues relating to web records, and is happy to do so upon request.

In addition to NARA's role to help agencies with guidance in managing all of their records, including web records, we are charged with the responsibility to provide for long-term preservation and access to those records selected for archival retention. The web site snapshot initiative outlined in this memorandum will ensure that we are able to document at least in part agency use of the Internet at the end of the Clinton Administration. As shown by the specifications for capturing and transferring this snapshot, NARA does not have the capability at this time to take or preserve all of the types of agency web records. NARA is currently working with the San Diego Supercomputer Center, InterPARES, and the international archival community in research projects addressing these issues.

LEWIS J. BELLARDO
Deputy Archivist of the United States

Attachment

NWM 05.2001 Memo to Agency Records Officers

Top of Page

Memos to Agency Records Officers
Records Management Home

National Archives and Records Administration home page
URL: http://www.nara.gov/records/cio-memo.html
webmaster@nara.gov
Last updated January 16, 2001

National Archives and Records Administration
Records Management

FAQ for this memo

January 12, 2001

NWM 05.2001

MEMORANDUM TO AGENCY RECORDS OFFICERS AND INFORMATION RESOURCE MANAGERS: Snapshot of Agency Public Web Sites

The National Archives and Records Administration (NARA) will be preserving a one-time snapshot of agency public web sites as they exist on or before January 20, 2001, as an archival record in the National Archives of the United States. The web site snapshot initiative outlined in this memorandum will ensure that we are able to document at least in part agency use of the Internet at the end of the Clinton Administration. To accomplish this, we ask that agencies take the following actions:

- Take a snapshot of your agency's public web site(s) following the instructions in the attachment to this memorandum.
- Within 60 days, send the snapshot and related documentation, also described in the attachment, to NARA.

I am sending this message to Chief Information Officers also. I urge you to work with your CIO to ensure that your agency's public web sites are included in the snapshot.

This immediate action does not address the larger and long-term records management issues relating to agency web sites that you and NARA face. Agencies are using the web more and doing more of their business on the web than ever before. Web sites include different types of records that are used in different ways, and document different agency activities of varying significance. Some web content parallels traditional paper records. In fact some agencies use the web to disseminate policy information, but publicize the fact that the authentic records of policy are maintained on paper. But some web sites involve records that are important evidence; some are without precedent. Web sites cannot be treated with a "one size fits all" solution and Federal agencies need comprehensive guidance on management of these records.

We will be working with agencies on the first phase of this guidance in the next few weeks. We will meet with records officers shortly to go over what they need to see in the guidance. Several agencies have submitted records schedules for their web sites and have received dispositions for them. NARA has also provided training and briefings on records management issues relating to web records, and is happy to do so upon request.

In addition to NARA's role to help agencies with guidance in managing all of their records, including web records, we are charged with the responsibility to provide for long-term preservation and access to those records selected for archival retention. The web site snapshot initiative outlined in this memorandum will ensure that we are able to document at least in part agency use of the Internet at the end of the Clinton Administration. As shown by the specifications for capturing and transferring this snapshot, NARA does not have the capability at this time to take or preserve all of the types of agency web records. NARA is currently working with the San Diego Supercomputer Center, InterPARES, and the international archival community in research projects addressing these issues.

LEWIS J. BELLARDO
Deputy Archivist of the United States

Attachment

1/12/01 Memorandum to Chief Information Officers

Top of Page
Memos to Agency Records Officers
Records Management Home

National Archives and Records Administration home page
URL: http://www.nara.gov /records/nwm05-01.html
webmaster@nara.gov
Last updated January 17, 2001

National Archives and Records Administration
Records Management

January 12, 2001

Attachment to NWM 05.2001

GUIDELINES TO AGENCIES ON PRESERVING A SNAPSHOT OF THEIR WEB SITES AT THE END OF THE CLINTON ADMINISTRATION

What does the one-time snapshot need to include?

A one-time snapshot of your agency's public Internet web site, taken on or before January 20, 2001. Your public web sites are those web sites maintained by the agency, directly or under contract, available for unrestricted viewing by the general public through the Internet. It does not include in-house or intranet web sites available only to employees of the agency. It should not include documents with information that is restricted in any way, such as security or privacy restricted information.

The snapshot needs to include all of the documents available to the public that are located on the agency's web server(s). In other words, the documents must be internal or contiguous to the web site. The snapshot should not include documents located on external servers to which the web site links. The agency should terminate those links.

Where the site or a page on the site is dynamic (i.e., the content exists in a database that serves the content through templates) take a snapshot of the template. Explain in the documentation that the file is a template that draws the information from a previously linked database (and give the title of the database) which is not included in the snapshot.

The snapshot should not be a back up of the system. It should be in a format that can be read on other platforms.

What is the procedure for sending the snapshot to NARA?

Agencies may send the snapshot on 3480-class tape cartridges, 9-track tapes, or CD-ROM. Media specifications are described in 36 CFR 1228.270 (http://www.nara.gov/nara/cfr/cfr1228l.html) and below:

- CD-ROMs should comply with ANSI/NISO/ISO 9660-1990—Volume and File Structure of CD-ROM for Information Exchange.
- Magnetic tape cartridges should be 18-track 3480-class tape cartridges recorded at 37,871 bpi that meet ANSI X3.180-1990, American National Standard: Magnetic Tape and Cartridge for Information Interchange. The data shall be blocked at no more than 32,760 bytes per block.
- 9-track tapes should be recorded at 1600 or 6250 bpi that meet ANSI X3.39-1986 or ANSI X3.54-1986.

For each web site, include in the package to NARA technical documentation adequate to identify, service and interpret the web site files. This should include:

- NARA Form 14097, Technical Description for Transfer of Electronic Records, and a completed NARA Form 14028, Information System Description Form, or their equivalents. If you need copies of these forms, contact NARA as provided below.
- Web site information content description documentation. This can be a copy of the site map or a list of files. The information should be sent in paper or ASCII form.
- Any information on the proprietary software and hardware used in constructing the web site that is needed to re-materialize the site.
- Information identifying the format of each file included in the snapshot, and where applicable, the

formats of any embedded objects.

Send the snapshot to:

National Archives and Records Administration
Electronic and Special Media Records Services Division (NWME)
8601 Adelphi Road
College Park MD 20740-6001

Who can I contact for further information?

For further information on transferring the files or taking the snapshot, and to request copies of the forms, please contact 301-713-6639 or CER@nara.gov.

How do I take the snapshot?

How you take the snapshot will depend on the level, depth, and availability of programming skills, and the format of the files. This methodology requires a low-level knowledge of file editing tools and web management skills (i.e. stream editing programs, web and platform server software configuration and administration, HTML, and JavaScript v1.2). The steps you need to take are summarized below.

1. Ensure that proper contingency procedures have been performed before proceeding (i.e., back up site).
2. Configure (e.g., mirror) your web site (or the portions to be transferred to NARA) in your developmental environment (e.g., another server).
3. Identify and calculate the number of external web site links.
4. Convert all external links by rerouting them back to original initiation points. Insert a notification message that the link has been terminated.
5. Identify all the files that need to be copied. See **What does the one-time snapshot need to include?** for information on the files that must be copied.
6. Copy the files to the appropriate media preserving the file and directory structures.

For further technical guidance, you may send an email to CER@nara.gov or contact 301-713-6639.

Return to Memo NWM 05.2001

National Archives and Records Administration
Records Management

January 12, 2001

Attachment to NWM 05.2001

GUIDELINES TO AGENCIES ON PRESERVING A SNAPSHOT OF THEIR WEB SITES AT THE END OF THE CLINTON ADMINISTRATION

What does the one-time snapshot need to include?

A one-time snapshot of your agency's public Internet web site, taken on or before January 20, 2001. Your public web sites are those web sites maintained by the agency, directly or under contract, available for unrestricted viewing by the general public through the Internet. It does not include in-house or intranet web sites available only to employees of the agency. It should not include documents with information that is restricted in any way, such as security or privacy restricted information.

The snapshot needs to include all of the documents available to the public that are located on the agency's web server(s). In other words, the documents must be internal or contiguous to the web site. The snapshot should not include documents located on external servers to which the web site links. The agency should terminate those links.

Where the site or a page on the site is dynamic (i.e., the content exists in a database that serves the content through templates) take a snapshot of the template. Explain in the documentation that the file is a template that draws the information from a previously linked database (and give the title of the database) which is not included in the snapshot.

The snapshot should not be a back up of the system. It should be in a format that can be read on other platforms.

What is the procedure for sending the snapshot to NARA?

Agencies may send the snapshot on 3480-class tape cartridges, 9-track tapes, or CD-ROM. Media specifications are described in 36 CFR 1228.270 (http://www.nara.gov/nara/cfr/cfr1228l.html) and below:

- CD-ROMs should comply with ANSI/NISO/ISO 9660-1990—Volume and File Structure of CD-ROM for Information Exchange.
- Magnetic tape cartridges should be 18-track 3480-class tape cartridges recorded at 37,871 bpi that meet ANSI X3.180-1990, American National Standard: Magnetic Tape and Cartridge for Information Interchange. The data shall be blocked at no more than 32,760 bytes per block.
- 9-track tapes should be recorded at 1600 or 6250 bpi that meet ANSI X3.39-1986 or ANSI X3.54-1986.

For each web site, include in the package to NARA technical documentation adequate to identify, service and interpret the web site files. This should include:

- NARA Form 14097, Technical Description for Transfer of Electronic Records, and a completed NARA Form 14028, Information System Description Form, or their equivalents. If you need copies of these forms, contact NARA as provided below.
- Web site information content description documentation. This can be a copy of the site map or a list of files. The information should be sent in paper or ASCII form.
- Any information on the proprietary software and hardware used in constructing the web site that is needed to re-materialize the site.
- Information identifying the format of each file included in the snapshot, and where applicable, the

formats of any embedded objects.

Send the snapshot to:

National Archives and Records Administration
Electronic and Special Media Records Services Division (NWME)
8601 Adelphi Road
College Park MD 20740-6001

Who can I contact for further information?

For further information on transferring the files or taking the snapshot, and to request copies of the forms, please contact 301-713-6639 or CER@nara.gov.

How do I take the snapshot?

How you take the snapshot will depend on the level, depth, and availability of programming skills, and the format of the files. This methodology requires a low-level knowledge of file editing tools and web management skills (i.e. stream editing programs, web and platform server software configuration and administration, HTML, and JavaScript v1.2). The steps you need to take are summarized below.

1. Ensure that proper contingency procedures have been performed before proceeding (i.e., back up site).
2. Configure (e.g., mirror) your web site (or the portions to be transferred to NARA) in your developmental environment (e.g., another server).
3. Identify and calculate the number of external web site links.
4. Convert all external links by rerouting them back to original initiation points. Insert a notification message that the link has been terminated.
5. Identify all the files that need to be copied. See **What does the one-time snapshot need to include?** for information on the files that must be copied.
6. Copy the files to the appropriate media preserving the file and directory structures.

For further technical guidance, you may send an email to CER@nara.gov or contact 301-713-6639.

Return to Memo NWM 05.2001

National Archives and Records Administration NARA

What's New on the RM Site

RM FAQs Index

CIO Link

Quick Links To:

Major Initiatives

Policy & Guidance

Communications

Opportunities for Comment

Locations

Resources

Training

Questions and Comments

Records Management Home

Federal Web Site Snapshot Information

- 1/12/01 Memo to CIOs, Snapshot of Agency Public Web Sites
- Attachment to Memorandum to CIOs, January 12, 2001, *Guidelines to Agencies on Preserving a Snapshot of Their Web Sites at the End of the Clinton Administration*
- NWM 05.2001 - Snapshot of agency public web sites
- FAQ for Records Officers, Snapshot of Agency Public Web Sites, 01/17/01
- FAQ on Technical Issues, Snapshot of Agency Public Web Sites, 01/18/01

FORMS

NA Form 14028 - Information System Description Form

- NA Form 14028 - (HTML version)
- NA Form 14028 Instructions - (HTML version)
- NA Form 14028 - (PDF version)
- NA Form 14028 - (Word97 version)
- NA Form 14028 Instructions - (Word97 version)

NA Form 14097 - Technical Description for Transfer of Electronic Records to the National Archives

- NA Form 14097 - (HTML version)
- NA Form 14097 Instructions - (HTML version)
- NA Form 14097 and Instructions - (Word97 version)

CIO Link - other information of interest to Chief Information Officers

Questions and comments

Top of Page
Records Management Home

SERIES_____ SUBSERIES_____

BOX_____ FOLDER_____

Information System Description Instructions

Definition

An *Information System* is the organized collection, processing, transmission, and dissemination of information in accordance with defined procedures. NARA's concern is with the government information in the system, that is, with information created, collected, processed, transmitted, disseminated, used, stored, and disposed of by the Federal Government. An electronic information system includes the inputs and outputs that are generated, as well as the information on electronic media. The system may contain budgetary, fiscal, social, economic, scientific-technical or program-related data and information, operated in support of agency programs and management responsibilities.

Explanations

1. The commonly used name and acronym of the system [e.g., Budget System, Grain Monitoring System]

2. The internal control number assigned to the system for reference, control, or cataloging purposes [e.g., Information System Inventory Number, ADP Plan control number]

3. What agency programs or missions does the system support?

4. What laws, directives, etc., authorize these programs?

5. Description has the following sections:

   a. *Purpose/Function:* The reasons for and the requirements met by the system.

   b. *Sources of Data:* The primary sources or providers of data to the system [e.g., broadcast license holders, corporations doing business in the U.S.]. Does this system receive information from other systems, either from within or outside your agency?

   c. *Information content:* The principal subject matter, data coverage, time span, geographic coverage, update cycle, whether the system saves superseded information, major characteristics of the system, and whether the system contains microdata or summary data.

   d. *Outputs:* The principal products of the system [e.g., reports, tables, charts, graphic displays, catalogs, correspondence], and an indication of the frequency of preparation. Is information from this system transferred to other systems?

6. Self-explanatory.

7. Self-explanatory.

8. Citations of previous NARA disposition jobs approving disposition of components [e.g.,

input forms. printouts, COM, output reports] of the system.

9. Self-explanatory.

NA FORM 14028 BACK (9-88)

Information System Description Form

1. SYSTEM TITLE	2. SYSTEM CONTROL NUMBER

3. AGENCY PROGRAM SUPPORTED BY SYSTEM	4. PROGRAM AUTHORITY

5. SYSTEM DESCRIPTION

5A. PURPOSE/FUNCTION OF SYSTEM

5B. SOURCE(S) OF DATA (Include inputs from other systems)

5C. INFORMATION CONTENT

5D. SYSTEM OUTPUTS (Include outputs from other systems)

6. NAME AND ADDRESS OF PRINCIPAL PROGRAM OFFICE SUPPORTED BY THE SYSTEM (Include room numbers)

7. AGENCY CONTACTS (Names, addresses, and phone numbers of system and program personnel who can provide additional information about the system and the program it supports.)

8. PREVIOUS DISPOSITION JOBS

9A. PREPARER'S NAME	**9B. OFFICE NAME AND ADDRESS**
9C. PHONE NUMBER	

SIGNATURE	**DATE**

TECHNICAL DESCRIPTION FOR TRANSFER OF ELECTRONIC RECORDS TO THE NATIONAL ARCHIVES

FILE IDENTIFICATION

OFFICIAL FILE TITLE, COMMONLY USED IDENTIFIER, AND/OR DESCRIPTIVE TITLE:	02. ACRONYM ASSIGNED TO FILE:
RESTRICTIONS ON ACCESS:	
04. TITLE/DESCRIPTION OF DOCUMENTATION PROVIDED:	05. FORMAT OF DOCUMENTATION: ☐ Paper ☐ Electronic Format (*Specify*) ☐ Other (Specify):
06. ELECTRONIC (Agency name and address):	07. IDENTIFY TECHNICAL CONTACT(S)

FILE CHARACTERISTICS

08. SHORT TITLE ON EXTERNAL LABEL	10. STORAGE MEDIA UNIT VOLUME SERIAL NUMBER:	12. DENSITY (CPI/BPI):
09. RETURN STORAGE MEDIA TO AGENCY AFTER ARCHIVAL PROCESSING: ☐ Yes ☐ No If Yes, Provide Address for Return (if different from item 06):	11. TYPE OF MEDIA PROVIDED: ☐ Open-Reel Magnetic Tape ☐ 3480-Class Tape Cartridge ☐ Other (Specify):	13. NUMBER OF TRACKS: ☐ 7 ☐ 9 ☐ 18 ☐ Other (Specify):
14. FILE ORGANIZATION ON STORAGE MEDIA ☐ One File on One Media Unit ☐ One File on Multiple Media Units ☐ Multiple Files on One Media Unit ☐ Multiple Files/Multiple Media Units	15. RECORDED LABEL (Internal Label) ☐ IBM OS ☐ IBM DOS ☐ ANSI X 3.27 Standard ☐ No Internal Labels ☐ Other (Specify and Describe)	16. CHARACTER SET ☐ ASCII ☐ EBCDIC ☐ BCD (7 track only) ☐ Binary ☐ Packed Decimal ☐ Other (Specify)
17. DATE FILE COPIED:		
18. INTERNAL FILE NAME/IDENTIFIER (aka Data Set Name):	19. SEQUENCE OF FILE ON STORAGE MEDIA UNIT: File of	20. NUMBER OF LOGICAL RECORDS (Blocking Factor):
21. RECORD TYPE ☐ Fixed Length (F) ☐ Fixed Length Blocked (FB) ☐ Other Than Fixed Length (Specify Format)	22. LENGTH OF LOGICAL RECORDS IN CHARACTERS OR BYTES: 23. LENGTH OF PHYSICAL RECORDS IN CHARACTERS OR BYTES:	24. NUMBER OF LOGICAL RECORDS: 25. NUMBER OF PHYSICAL RECORDS (Blocks):
26. AGENCY COMMENTS:		
27. FORM PREPARED BY: Name:	Phone:	28. DATE FORM COMPLETED:

INSTRUCTIONS ON OTHER SIDE

INSTRUCTIONS

GENERAL:

The purpose of this form is to facilitate the transfer of electronic records with continuing or enduring value to the National Archives. A packet containing additional information on the transfer of electronic records is available from the National Archives. This form is not intended to take the place of other required documentation. FIPS PUB 20 contains a discussion of file documentation. Technical information describing each file is required by 36 CFR, but NA Form 14097 is optional. If there is an alternative reporting format that provides all of this required information, substitute the report for NA Form 14097. Include the required information as an attachment to the SF 258. A separate form should be completed for each file. If multiple files have very similar technical specifications, one form with an attachment that specifically identifies all of the files covered by the form may be used.

IDENTIFICATION SECTIONS

01. **Official Title, Commonly Used Identifier, and/or Descriptive Title.** Enter the name by which the agency identifies the file. Consider how the title would appear in a bibliographic entry. If there is no official title, provide a descriptive title.

02. **Acronym Assigned to File.** Enter the commonly used abbreviation or acronym as assigned by the agency. Often, the acronym as assigned by the agency. Often, the acronym will be used on the external (gummed) label of the storage media unit.

03. **Restrictions on Access.** Specify any restrictions that apply to this file - cite FOIA exemption, and, if b (3), cite statute, indicate specific columns of types of records in the file that are affected; specify length of restriction on access and method of determining the date when restrictions end. If there are no applicable restrictions on access, please indicate.

04. **Title/Description of Documentation Provided.** Documentation is required for all transfers of electronic records to the National Archives. Enter the title or description of the documentation provided by the agency for the file. Guidelines are available on the source and content of documentation. If any documentation is available in electronic form, include it in the transfer.

05. **Format of Documentation Provide.** Mark all boxes that apply to the transfer with an "X." If "Electronic Format" is checked, include a technical description form for each documentation file in electronic format. If "Other" is checked, be as specific as possible in describing the documentation transferred.

06. **Electronic Records Submitted by.** Enter the name and address of the agency that is responsible for the transfer.

07. **Identify Technical Contact(s).** Identify the person who will respond to technical questions about the records if they arise during archival processing.

FILE CHARACTERISTICS

08. **Short Title on External Label of Storage Media Unit.** Enter the short title that appears on the external (gummed) label of the storage media unit(s).

09. **Return Storage Media to Agency After Archival Processing.** The National Archives returns the reels or cartridges included in the transfer to the agency when two preservation copies have been created. Check "No" to indicate that the storage media should not be returned or provide an address for return shipment. If yes, Provide Address for Return (if different from Item 06).

10. **Storage Media Unit Volume Serial Number.** Enter the volume serial number which uniquely identifies this tape/tape cartridge. If the file is recorded on multiple

volumes, enter the first volume serial number in this item and list others, in sequence, in item 26.

11. **Type of Media Provided.** To comply with the transfer standard identified in 36 CFR, a storage media unit should be an open-reel magnetic tape or 3480 tape cartridge. Enter an "X" in the appropriate box. If "Other" is checked, contact the National Archives prior to transfer and provide a specific identification of the storage media used for transfer.

12. **Density (CPI/BPI).** Enter an "X" in the appropriate box to indicate characters or bytes per inch.

13. **Number** of Tracks. Enter an "X" in the appropriate box. For 7 track tapes, indicate whether the parity is odd or even in Item 26.

14. **File Organization on Storage Media.** If a single file on a single storage media unit is described, check "One File on One Media Unit." If a multi-volume file is described, check "One File on Multiple Media Units." If more than one file is on the storage media unit, check "Multiple Files on One Media Unit." If the transfer includes multiple files on multiple files on One media units, check "Multiple Files on Multiple Media Units."

15. **Recorded Label (Internal Label).** Enter an "X" in the appropriate box. If "Other" is checked, provide a specific description of the internal labels in Item 26.

16. **Character Set.** Enter an "X" in the appropriate box. If "Binary" or "Packed Decimal" is checked, indicate the characters (bytes) that are stored in those formats in Item 26.

17. **Date File Copied.** Enter the date that appears on internal labels, if the files are labelled. This is the date the records were copied onto the storage media unit(s).

18. **Internal File Name/Identifier (aka Data Set Name).** If there is a recorded label, enter the file identification (e.g., Data Set Name) used in the label. File Names in IBM standard labels should follow IBM DSN naming conventions.

19. **Sequence of File on Storage Media Unit.** If this is the only file on the storage media unit, enter a "1." If the storage media unit contains multiple files, enter the file's position number on the storage media unit. See Item 14.

20. **Number of Logical Records per Block (Blocking Factor).** If "Fixed Length" is checked in Item 21, enter the blocking factor.

21. **Record Type.** Enter an "X" in the appropriate box. If "Other Than Fixed Length" is checked, a specific description of the format is very important, especially if the documentation provided does not contain a precise description. Use Item 26 to describe "Other" formats.

22. **Length of Logical Records in Characters or Bytes.** Enter the logical record length in bytes or characters. If the record is not fixed length, enter the maximum size record, and describe techniques used to control and indicate size in Item 26.

23. **Length of Physical Records in Characters or Bytes.** Enter the physical record length in bytes or characters. If the record is not fixed length, enter the maximum size record, and describe techniques used to control and indicate size in Item 26.

24. **Number of Logical Records.** Enter the number of logical records in the file. This number is usually supplied by the program which created the file. Labels are not included in this count. If the last block is padded with blank records to fill out the block, please provide a total record count and a count of records that contain information.

25. **Number of Physical Records (Blocks).** Enter the number of

physical records (blocks). Labels are not included in this count.

26. **Comments.** Provide additional information as necessary or desired. If the information pertains to another item on the form, identify the relevant item number. Files transferred to the National Archives should not be software-dependent in accordance with 36 CFR. If there are any exceptions to this, identify the release and/or level of any software required to read the file. Provide attachments if the information required will not fit in Item 26.

27. **Form Prepared By.** Enter the name and phone number (including area code) of the individual who prepared this form.

28. **Date Form Completed (YY/MM/DD).** Enter date this form was prepared.

http://search.npr.org/cf/cmn/cmnpd01fm.cfm?PrgDate=9%2F2%2F1999&PrgID=2

Internet -- It's a birthday party for the Internet! Today at UCLA many of the scientists who made it happen, are celebrating their 30-year-old creation. Robert Smith reminds us that in the "summer of '69" along with "Abbey Road," men on the moon and Woodstock, the first computer -- or node -- on the Internet went live at UCLA. (5:30)

Laura O'Hara
9/3/99

SERIES _____ SUBSERIES _____

BOX 16 FOLDER 17

Received: from SMTP.SLAC.Stanford.EDU ([134.79.18.80]) by atreus.slac.
stanford.edu with SMTP (Microsoft Exchange Internet Mail Service Versi
on 5.5.2653.13)
id NP91FX4Q; Thu, 5 Jul 2001 14:07:59 -0700
Received: from CONVERSION-DAEMON by smtp.slac.stanford.edu
(PMDF V5.2-33 #37476) id <0GG000701QPCVV@smtp.slac.stanford.edu> for
jmdeken@exchange.slac.stanford.edu (ORCPT rfc822;jmdeken@slac.stanfor
d.edu)
; Thu, 5 Jul 2001 14:08:01 -0700 (PDT)
Received: from directory-daemon by smtp.slac.stanford.edu (PMDF V5.2-3
3 #37476)
id <0GG000701QPCVU@smtp.slac.stanford.edu> for
jmdeken@exchange.slac.stanford.edu (ORCPT rfc822;jmdeken@slac.stanfor
d.edu)
; Thu, 05 Jul 2001 14:08:00 -0700 (PDT)
Received: from postal.sdsc.edu (postal.sdsc.edu [132.249.20.114])
by smtp.slac.stanford.edu (PMDF V5.2-33 #37476)
with ESMTP id <0GG0002JSQPCTK@smtp.slac.stanford.edu> for
jmdeken@SLAC.Stanford.EDU; Thu, 05 Jul 2001 14:08:00 -0700 (PDT)
Received: from ptlc.sdsc.edu (ptlc.sdsc.edu [132.249.24.111])
by postal.sdsc.edu (8.9.3/8.9.3/SDSCserver-17) with ESMTP id OAA01975
; Thu,
05 Jul 2001 14:07:59 -0700 (PDT)
Date: Thu, 05 Jul 2001 14:07:59 -0700 (PDT)
From: Joseph Deken <dekenj@SDSC.EDU>
Subject: Ginsparg & archive to Cornell
To: jwooley@ucsd.edu, jmdeken@SLAC.Stanford.EDU
Message-id: <Pine.GSO.4.30.0107051405450.21077-100000@ptlc.sdsc.edu>
MIME-version: 1.0
Content-type: TEXT/PLAIN; charset=X-UNKNOWN
Content-transfer-encoding: QUOTED-PRINTABLE

This is as interesting for what it says about Cornell as it does abou=
t
LANL.

05 July 2001

Nature 412, 3 - 4 (2001) =A9 Macmillan Publishers Ltd.

Los Alamos loses physics archive as preprint pioneer heads east

DECLAN BUTLER

2000 UC

Low vantage point: the loss of the preprint server is a blow for the

lab on the New Mexico mesa.

The Los Alamos preprint server, which has established itself as physicists' favourite place for early circulation of their research, is leaving the New Mexico laboratory to set up shop at Cornell University in New York state.

Paul Ginsparg, who founded the server =97 now known as arXiv 10 years ago, is leaving the Los Alamos National Laboratory (LANL) to take up = a
faculty position at Cornell, and the server will move with him. Cornell plans to expand arXiv's reach into other disciplines, and to use it as a test bed for research into digital libraries.

Ginsparg says growing dissatisfaction with LANL is a major reason for his departure, citing a lack of enthusiasm for the archive among senior staff. Only his former group leader Geoffrey West and library director Rick Luce gave the archive strong support, he says. He adds that the nuclear-weapons laboratory has been shifting its support towards large groups at the expense of individual investigators, and is suffering from declining morale in the wake of recent security scandals.

Los Alamos experienced a painful security clamp-down after Wen Ho Lee= ,
a Taiwanese-born engineer at the laboratory, was arrested two years ago on espionage charges and then released (see next article). The loss of the prestigious server delivers another blow to the laboratory's standing in the scientific community.

Paul Ginsparg: set to move to Cornell.

William Press, deputy director of the laboratory, says: "We're sorry to see Paul go, but Cornell has created a very unique opportunity for him. We are very proud to have been the incubator of this revolution in scientific publishing." He adds that senior laboratory staff have strongly supported the archive activity, but admits that it was sometimes "a struggle to see where it would fit in" with the laboratory's other activities.

The archive currently receives around $300,000 in annual funding from the National Science Foundation, the Department of Energy, which runs the lab, and LANL itself.

Ginsparg says that consultation with the archive's advisory board, funding agencies and the American Physical Society, produced a consensus that the operation would enjoy more secure funding and stronger intellectual support at a university than at LANL.

But for Ginsparg, the last straw was his recent salary review, which, he says, described him as "a strictly average performer by overall la=
b
standards; with no particular computer skills contributing to lab programs; easily replaced, and moreover overpaid, according to an external market survey".

LANL officials declined to comment on Ginsparg's case, but said that some recent salary increases at the laboratory have been available only to certain combinations of programmes and individual skills.

Peter Lepage, chair of Cornell's physics department, notes wryly of the LANL assessment: "Evidently their form didn't have a box for: 'completely transformed the nature and reach of scientific informatio=
n
in physics and other fields'."

--end--

SERIES _____ SUBSERIES _____
BOX _____ FOLDER _____

NEWS: An angry sky: Surprise lightning show crackles over the Bay Area/2B

San Jose Mercury News

PENINSULA EDITION
35 CENTS

Serving Northern California Since 1851

THURSDAY
... SEPTEMBER 9, 1999

WWW.MERCURYCENTER.COM

Finally, the Internet has a birth date

DAVID PLOTNIKOFF
MODEM DRIVER

THIRTY years have not dimmed our collective memory of Neil Armstrong's walk on the moon. How is it that the *other* giant leap of 1969 — the birth of the Internet — could be so obscured by the passage of time?

For decades, none of the researchers who built and tested that first Net connection have been able to recall the exact date of the transmission. Although the Net has been the subject of innumerable articles, books and documentaries in recent years, none has been able to establish a birth date for the mother of all networks.

Ultimately, it was the death of one of the Net's founding fathers that brought the date to light. Based on that recently rediscovered documentation and new recollections by individuals on both ends of that first communication, I'm pleased to say that for the first time the Internet has a real birthday — Oct. 29, 1969.

Actually, it has *two* birthdays — one that's passed and the one next month. And therein lies a bit of a political problem. Given the propensity for endless debate among historians and Net-geeks of long standing, it's quite possible scholars could be slugging it out over the date right up until the Net turns 40.

There is a plausible explanation for how such a key milestone could be effectively forgotten for so long: In the race to get the network up and running, that one communication simply wasn't considered terribly important. "The truth is, nobody paid much attention to it," said Doug Engelbart, the computer vision-

See NET BIRTHDAY, Back Page

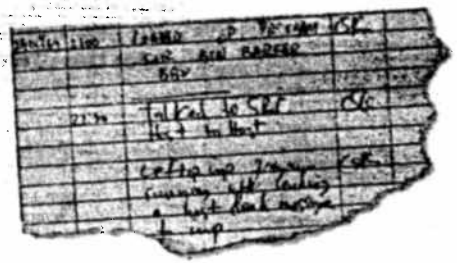

Charley Kline's logbook documents logging in to a Stanford computer from UCLA at 10:30 p.m. on Oct. 29, 1969.

MERCURY NEWS FILE PHOTOGRAPH

In 1969, Doug Englebart ran a computer lab at Stanford Research Institute in Menlo Park, which was the receiving end of the Internet's first remote log-in.

ASSOCIATED PRESS

Leonard Kleinrock says his watch has more processing power than the IMP, the very first node of the Internet.

Finally, Internet has a birth date: Oct. 29, 1969

■ NET BIRTHDAY

from Page 1A

ary who oversaw one end of the early network back in '69. "It was more like 'Well, OK, they brought in the motor so we can get our vehicle running.' It was just part of getting the plumbing to work."

The basic outline of the Net's Genesis story is clear: The first node of the Pentagon-funded Advanced Research Projects Agency network (ARPAnet) was installed in the lab of a UCLA computer-science professor, Leonard Kleinrock, on Labor Day weekend of 1969. If you want to get technical about it, what happened is this: A specialized computer called an "Interface Message Processor" (IMP) was installed in the lab. That machine, a battleship gray Honeywell Model 516 minicomputer (the size of a refrigerator) would act as intermediary between the lab's mainframe, a Scientific Data Systems Sigma-7, and the nascent network. On Sept. 2, the UCLA team succeeded in moving data back and forth from the Sigma-7 to the IMP, which stood 15 feet away.

DAVID PLOTNIKOFF
MODEM DRIVER

What followed is slightly less clear. Approximately one month later, the second IMP was installed in Engelbart's lab at Stanford Research Institute in Menlo Park. IMP No. 2 was connected to Engelbart's Scientific Data Systems 940 mainframe.

The two labs — in Los Angeles and Menlo Park — were linked by a leased 50-kilobit-per-second dedicated phone line that enabled both voice and data traffic between the two labs. As many of the researchers who were present at both ends recalled, the first data communication between the two labs' host computers occurred "sometime in October." Five years ago, as these pioneers prepared to gather in Boston for a 25-year reunion, I spoke to a half-dozen of them and came away with the distinct impression that "sometime in October" would be as close as we'd ever get to the truth.

Prior to that 1994 anniversary, Jon Postel, who went from UCLA to become the Net's ultimate keeper of names and numbers, told me: "You have to understand the context. Here we were doing this little research project that involved a few computer scientists. There wasn't any thought that we were exploring this brave new world or inventing it."

Today the Net is a magnitude larger than it was just five years ago, and the lack of a birthday can't be shrugged off or finessed. The wired world is now such an overwhelming part of mass culture that it demands an anniversary date. This year, Kleinrock and UCLA's public relations machine were only too happy to oblige. A few weeks back, much to the dismay of some observers who were not affiliated with that university, UCLA announced unilaterally that Sept. 2 would be the Internet's official birthday.

The problem with this declaration is simple: UCLA's IMP was talking to UCLA's own computer, but it was not, strictly speaking, a test of the network. It was the packet-switching equivalent of a single tin can and a string. The Internet would not exist until there was someone on the other end of the line.

Sometime in the weeks after IMP No. 2 arrived at SRI, Charley Kline, a grad student in Kleinrock's lab at UCLA, made the first attempt to log into SRI's machine via the network. For this particular test, his machine at UCLA tried to connect with SRI and execute commands on the SRI machine as if he were right there in the lab. His first attempt to log in (there was no "first message" or e-mail per se, since e-mail as we know it had yet to be invented) caused SRI's computer to crash.

On Tuesday, Kline, now a senior technologist at Cisco Systems Inc. in San Jose, related to me how he found himself on the initiating end of the historic connection. "I was the operating-systems programming kid at that time. So I was the guy who had to write the code to make the IMP talk to the host," said Kline, 51. "I ended up writing something similar to a telnet program (a program that allows a user to log in remotely on a distant system) and trying to log into SRI's computer, making it look as if we were a dumb terminal."

As Kline recalls, a day or two after his first failed attempt, he tried again — and it worked.

The date of that remote log-in was promptly forgotten in the rush to hook up additional nodes.

Last October, Postel, who'd worked alongside Kline for many years, died after complications from surgery. Kline attended three memorial services in the Los Angeles area for Postel. After the UCLA service, Kleinrock e-mailed his eulogy to attendees — along with some historic images he'd pulled from his archives. One of the images, a page from a log book for IMP No. 1, caught Kline's eye. He saw, in his own handwriting, the date "29OCT69" and the terse notation that at 10:30 p.m. he "talked to SRI host to host." (Kline notes now that he may have had a previous successful attempt to log in. He's just not sure. The Oct. 29 connection was the one he knows he documented.)

Judy Lin-Eftekhar, a spokeswoman for UCLA's school of engineering, said Wednesday that the university is "aware of the controversy" over the official birthday. "We at UCLA are saying it's Sept. 2 because because that's when the first network switch was put into place," she said. "That technology basically enabled the Internet to follow. That's our claim, but we're aware others may see things differently."

Is this an utterly trivial issue? Many of the pioneers present at both institutions 30 years ago seem to think so. I disagree. We must find some mechanism, such as a birthday, to look back on whence we came. Today it sounds vaguely Marxist, given the gold-rush mindset surrounding the Net, but guys like Engelbart, Kleinrock and their respective teams set out to build a platform for collaborative work — a public resource.

Their core values — cooperation, governance by consensus, open standards — were what enabled the Net to scale up a million-fold without breaking. We should not wait for the next round of eulogies. We should honor these people and the values that steered them now. Anyone whose life has been touched for the better by the Net has some moral obligation to say thank you. Why not Oct. 29?

LEN LAHMAN — MERCURY NEWS

Charley Kline was a UCLA graduate student in October 1969 when he successfully logged in to a computer in Menlo Park. Now he works for Cisco.

David Plotnikoff writes about the wired life for the San Jose Mercury News. Contact him at plotnikoff@sjmercury.com or (408) 920-5867. On the Web, dial
www.sjmercury.com/columnists/plotnikoff

NEWS: An angry sky: Surprise lightning show crackles over the Bay Area/2B

San Jose Mercury News

PENINSULA EDITION
35 CENTS

Serving Northern California Since 1851

THURSDAY
SEPTEMBER 9, 1999

WWW.MERCURYCENTER.COM

Finally, the Internet has a birth date

DAVID PLOTNIKOFF
MODEM DRIVER

THIRTY years have not dimmed our collective memory of Neil Armstrong's walk on the moon. How is it that the *other* giant leap of 1969 — the birth of the Internet — could be so obscured by the passage of time?

For decades, none of the researchers who built and tested that first Net connection have been able to recall the exact date of the transmission. Although the Net has been the subject of innumerable articles, books and documentaries in recent years, none has been able to establish a birth date for the mother of all networks.

Ultimately, it was the death of one of the Net's founding fathers that brought the date to light. Based on that recently rediscovered documentation and new recollections by individuals on both ends of that first communication, I'm pleased to say that for the first time the Internet has a real birthday — Oct. 29, 1969.

Actually, it has *two* birthdays — one that's passed and the one next month. And therein lies a bit of a political problem. Given the propensity for endless debate among historians and Net-geeks of long standing, it's quite possible scholars could be slugging it out over the date right up until the Net turns 40.

There is a plausible explanation for how such a key milestone could be effectively forgotten for so long: In the race to get the network up and running, that one communication simply wasn't considered terribly important. "The truth is, nobody paid much attention to it," said Doug Engelbart, the computer vision-

See NET BIRTHDAY, Back Page

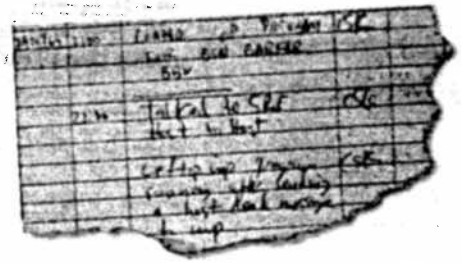

Charley Kline's logbook documents logging in to a Stanford computer from UCLA at 10:30 p.m. on Oct. 29, 1969.

MERCURY NEWS FILE PHOTOGRAPH

In 1969, Doug Englebart ran a computer lab at Stanford Research Institute in Menlo Park, which was the receiving end of the Internet's first remote log-in.

Finally, Internet has a birth date: Oct. 29, 1969

ASSOCIATED PRESS

Leonard Kleinrock says his watch has more processing power than the IMP, the very first node of the Internet.

■ NET BIRTHDAY

from Page 1A

ary who oversaw one end of the early network back in '69. "It was more like 'Well, OK, they brought in the motor so we can get our vehicle running.' It was just part of getting the plumbing to work."

The basic outline of the Net's Genesis story is clear: The first node of the Pentagon-funded Advanced Research Projects Agency network (ARPAnet) was installed in the lab of a UCLA computer-science professor, Leonard Kleinrock, on Labor Day weekend of 1969. If you want to get technical about it, what happened is this: A specialized computer called an "Interface Message Processor" (IMP) was installed in the lab. That machine, a battleship gray Honeywell Model 516 mincomputer (the size of a refrigerator) would act as intermediary between the lab's mainframe, a Scientific Data Systems Sigma-7, and the nascent network. On Sept. 2, the UCLA team succeeded in moving data back and forth from the Sigma-7 to the IMP, which stood 15 feet away.

DAVID PLOTNIKOFF
MODEM DRIVER

What followed is slightly less clear. Approximately one month later, the second IMP was installed in Engelbart's lab at Stanford Research Institute in Menlo Park. IMP No. 2 was connected to Engelbart's Scientific Data Systems 940 mainframe.

The two labs — in Los Angeles and Menlo Park — were linked by a leased 50-kilobit-per-second dedicated phone line that enabled both voice and data traffic between the two labs. As many of the researchers who were present at both ends recalled, the first data communication between the two labs' host computers occurred "sometime in October." Five years ago, as these pioneers prepared to gather in Boston for a 25-year reunion, I spoke to a half-dozen of them and came away with the distinct impression that "sometime in October" would be as close as we'd ever get to the truth.

Prior to that 1994 anniversary, Jon Postel, who went from UCLA to become the Net's ultimate keeper of names and numbers, told me: "You have to understand the context. Here we were doing this little research project that involved a few computer scientists. There wasn't any thought that we were exploring this brave new world or inventing it."

Today the Net is a magnitude larger than it was just five years ago, and the lack of a birthday can't be shrugged off or finessed. The wired world is now such an overwhelming part of mass culture that it demands an anniversary date. This year, Kleinrock and UCLA's public relations machine were only too happy to oblige. A few weeks back, much to the dismay of some observers who were not affiliated with that university, UCLA announced unilaterally that Sept. 2 would be the Internet's official birthday.

The problem with this declaration is simple: UCLA's IMP was talking to UCLA's own computer, but it was not, strictly speaking, a test of the network. It was the packet-switching equivalent of a single tin can and a string. The Internet would not exist until there was someone on the other end of the line.

Sometime in the weeks after IMP No. 2 arrived at SRI, Charley Kline, a grad student in Kleinrock's lab at UCLA, made the first attempt to log into SRI's machine via the network. For this particular test, his machine at UCLA tried to connect with SRI and execute commands on the SRI machine as if he were right there in the lab. His first attempt to log in (there was no "first message" or e-mail per se, since e-mail as we know it had yet to be invented) caused SRI's computer to crash.

On Tuesday, Kline, now a senior technologist at Cisco Systems Inc. in San Jose, related to me how he found himself on the initiating end of the historic connection. "I was the operating-systems programming kid at that time. So I was the guy who had to write the code to make the IMP talk to the host," said Kline, 51. "I ended up writing something similar to a telnet program (a program that allows a user to log in remotely on a distant system) and trying to log into SRI's computer, making it look as if we were a dumb terminal."

As Kline recalls, a day or two after his first failed attempt he tried again — and it worked. The date of that remote log-in was promptly forgotten in the rush to hook up additional nodes.

Last October, Postel, who'd worked alongside Kline for many years, died after complications from surgery. Kline attended three memorial services in the Los Angeles area for Postel. After the UCLA service, Kleinrock e-mailed his eulogy to attendees — along with some historic images he'd pulled from his archives. One of the images, a page from a log book for IMP No. 1, caught Kline's eye. He saw, in his own handwriting, the date "29OCT69" and the terse notation that at 10:30 p.m. he "talked to SRI host to host." (Kline notes now that he may have had a previous successful attempt to log in. He's just not sure. The Oct. 29 connection was the one he knows he documented.)

LEN LAHMAN — MERCURY NEWS

Charley Kline was a UCLA graduate student in October 1969 when he successfully logged in to a computer in Menlo Park. Now he works for Cisco.

Judy Lin-Eftekhar, a spokeswoman for UCLA's school of engineering, said Wednesday that the university is "aware of the controversy" over the official birthday. "We at UCLA are saying it's Sept. 2 because because that's when the first network switch was put into place," she said. "That technology basically enabled the Internet to follow. That's our claim, but we're aware others may see things differently."

Is this an utterly trivial issue? Many of the pioneers present at both institutions 30 years ago seem to think so. I disagree. We must find some mechanism, such as a birthday, to look back on whence we came. Today it sounds vaguely Marxist, given the gold-rush mindset surrounding the Net, but guys like Engelbart, Kleinrock and their respective teams set out to build a platform for collaborative work — a public resource.

Their core values — cooperation, governance by consensus, open standards — were what enabled the Net to scale up a million-fold without breaking. We should not wait for the next round of eulogies. We should honor these people and the values that steered them now. Anyone whose life has been touched for the better by the Net has some moral obligation to say thank you. Why not Oct. 29?

David Plotnikoff writes about the wired life for the San Jose Mercury News. Contact him at plotnikoff@sjmercury.com or (408) 920-5867. On the Web, dial
www.sjmercury.com/columnists/plotnikoff

USA TODAY Tech Report

09/01/99- Updated 10:34 AM ET

Who really invented the Net?

4 'forefathers' differ on who deserves credit

By Bruce Haring, USA TODAY

The date was Oct. 29, 1969. The laboratory of professor Leonard Kleinrock at the University of California at Los Angeles was about to give birth to the Internet.

Kleinrock's research assistant, Charley Kline, was wearing a communications headset as he sat at a terminal hooked to an interface message processor, the first computer network switch. On Kleinrock's command, Kline spoke with a colleague at Stanford University in the Bay Area,

What lies ahead?
The Net's founders predict its future

"Nomadic computing, providing access while you're on the road so that the Internet services you see when you're

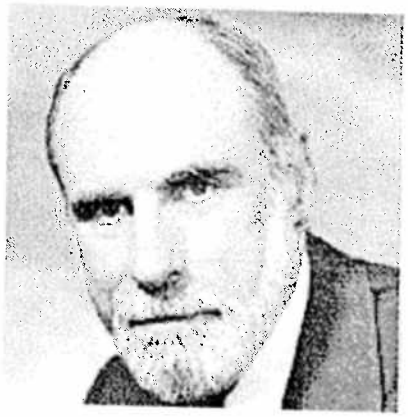

"Radio-based links into the Net will be very typical. If you have a question, you'll whip

Inside Tech
FAQ/Tips
Web Column
Hot Sites
Tech News
Tech Reviews
Answer Desk
Game Zone
Shareware Shelf
Web Potholes
Web Resources
Tech Front

Hardware
Accessories
Software

Search
Newspaper Archives
Our site

Resources
Index
Search
Feedback
What's hot
About us
Jobs at USA TODAY

the other end of the long-distance, high-speed line.

"All we wanted to do was (type the words) 'log in,' " recalls Kleinrock, whose groundbreaking research in network data-packet switching led to the historic moment.

Kline typed the "L." Stanford acknowledged that the letter had been received.

Kline typed the "O." Also received.

Kline typed the "G."

The system crashed.

The first message sent across the Net turned out to be quite appropriate -- either the biblical word "lo" or, phonetically, a version of "hello," Kleinrock says.

services you see when you're someplace else are no different than what you have back in your office."
--Leonard Kleinrock

"Many sites in the research community will have access at gigabyte speed to the Internet. You'll see the increasing introduction of wireless access, so people don't have to feel tethered to the Net. And we're going to see increasing content."
--Robert Kahn

out your Palm Pilot with a radio link and go on the Net and pull the data out."
--Vinton Cerf

"The Internet will become the pervasive network for the world's telecom traffic. Voice and video will transfer over to it in the next five to 10 years. Clearly, you're going to have video on demand, radio or TV, that can have millions of different sources or special subjects that (small numbers) care about."
--Lawrence Roberts

An hour later, they got the system up and running and completed the experiment as planned. The Net was born -- at least, some contend it was.

In the 30 years since, millions of messages and files have traveled across the worldwide computer network that sprang from that humble beginning. Kleinrock, with Robert Kahn, Vinton Cerf and Lawrence G. Roberts -- the three other men considered key to what today is known as the Internet -- will speak Thursday at a UCLA symposium devoted to the medium's birth, its current state and its future.

Kleinrock is still at UCLA and also is head of a company called Nomadix, devoted to

computer use beyond the desktop. Kahn is president of the Corporation for National Research Initiatives, a not-for-profit organization devoted to research for the nation's information infrastructure. Cerf is MCI's senior vice president for Internet architecture and technology, and Roberts is president and CEO of Packetcom.com, a communications-switching research firm.

Though they share a crucial role in shaping electronic communications, the four will not be on the same screen regarding certain Net issues, including who can properly lay claim to "fathering" the Internet.

Kleinrock's experiment made his computer the first node of the Advanced Research Projects Agency's ARPANET, the world's first major packet-data transmission network. Conceived in the United States' scientific scramble after the Soviet Union's launch of the Sputnik satellite, ARPANET was a single network linking some of the nation's research computers.

As time went on, other networks were built. Many contend that the ability to link from any computer into any network (research that began in the 1970s and was fully implemented in 1983) is what forms the Net we know today.

For those who subscribe to that Internet creation theory, Cerf and Roberts, who co-designed the architecture for the collective network system and its basic communications protocols, TCP/IP, are the true fathers of the Internet.

Related stories:

- Debates persist over Net's birthday
- Humble beginnings: Net turns 30

Because the issue is clearly a matter of perspective, the founding fathers have reached something of an uneasy standoff on the issue.

"I think that the old arguments that will come up at the (UCLA) conference and have come up over and over is everybody is claiming responsibility for everything at this point," says Roberts, who was the designer and developer of ARPANET.

But one thing all agree on is that the Internet was not conceived as a fail-safe communications tool in case of nuclear war, a much-promulgated myth over the years. The Rand Research Institute was developing a study shortly after ARPANET's birth that has been confused with the research-oriented ARPANET and subsequent developments.

Nuclear war "wasn't the reason we did anything," Roberts says. "That story is just wrong."

Despite such myths and disputes, the four "forefathers" share a pride in the commerce, entertainment and other opportunities that have sprung from their research.

And one regret.

"I didn't get to play with this stuff until I was 28," laments Cerf, echoing his colleagues.

"I often envy the kids who are 8 years old and hacking the Net. They got to do it 20 years before I did. I had to go off and invent it."

---

http://www.wired.com/news/culture/0,1284,32073,00.html

Interview with the Web's Creator

by Chris Oakes

3:00 a.m. 23.Oct.1999 PDT

Tim Berners-Lee has finally spoken. His new book, Weaving the Web, chronicles how the Web really happened and where its creator thinks it should go from here.

In 1989, at the European Particle Physics Laboratory in Geneva, Switzerland, Berners-Lee first proposed a "global hypertext project" to be known as the World Wide Web. He wanted researchers like himself to be able to easily and automatically combine their knowledge in a Web of hypertext documents.

Berners-Lee wrote the software that would eventually open the Web browser and provide its underlying protocols: the hypertext markup language, or HTML, and the hypertext transfer protocol, the http:// that precedes the zillions of Web addresses -- Uniform Resource Locators, or URLs -- that are now as ubiquitous as ZIP codes.

Such topics may be as enticing to lay readers as fingernails on a chalkboard. But when Berners-Lee starts talking them up, things get interesting. Such topics are fundamental to his once and future vision for the Web as a body of living intelligence, a place where any piece of information can be linked to any other piece of information. Berners-Lee is still dealing with the fact that that vision is far from realized.

He remains director of the Web's standards-bearer, the World Wide Web Consortium, which he founded shortly after arriving at MIT in 1994. It may have seemed a simple decision to go nonprofit, but his objective was to ensure that the W3C remained universally accessible. As he told Wired News, the job was not as easy as it looked and, even today, it's a challenge.

Wired News: Are you comfortable being referred to as the inventor of the World Wide Web?

Tim Berners-Lee: I don't mind being, in the public context, referred to as the inventor of the World Wide Web. What I like is that image to be separate from private life, because celebrity damages private life.

Interview with the Web's Creator
by Chris Oakes

3:00 a.m. 23.Oct.1999 PDT
Tim Berners-Lee has finally spoken. His new book, Weaving the Web, chronicles how the Web really happened and where its creator thinks it should go from here.

In 1989, at the European Particle Physics Laboratory in Geneva, Switzerland, Berners-Lee first proposed a "global hypertext project" to be known as the World Wide Web. He wanted researchers like himself to be able to easily and automatically combine their knowledge in a Web of hypertext documents.

Listen to the audio version

Berners-Lee wrote the software that would eventually open the Web browser and provide its underlying protocols: the hypertext markup language, or HTML, and the hypertext transfer protocol, the http:// that precedes the zillions of Web addresses -- Uniform Resource Locators, or URLs -- that are now as ubiquitous as ZIP codes.

Such topics may be as enticing to lay readers as fingernails on a chalkboard. But when Berners-Lee starts talking them up, things get interesting. Such topics are fundamental to his once and future vision for the Web as a body of living intelligence, a place where any piece of information can be linked to any other piece of information. Berners-Lee is still dealing with the fact that that vision is far from realized.

He remains director of the Web's

standards-bearer, the World Wide Web Consortium, which he founded shortly after arriving at MIT in 1994. It may have seemed a simple decision to go nonprofit, but his objective was to ensure that the W3C remained universally accessible. As he told Wired News, the job was not as easy as it looked and, even today, it's a challenge.

Wired News: Are you comfortable being referred to as the inventor of the World Wide Web?

Tim Berners-Lee: I don't mind being, in the public context, referred to as the inventor of the World Wide Web. What I like is that image to be separate from private life, because celebrity damages private life.

WN: Is it an appropriate label though? HTTP, HTML, URLs are all your inventions, right? So it seems that the label does apply.

Berners-Lee: Pretty much. I basically wrote the code and the specs and documentation for how the client and server talked to each other.

WN: You had the Internet itself to inspire you of course in some of the protocols that were already there.

Berners-Lee: Oh, I had a whole lot. In fact, a lot of the design decisions were not only using that experience, but also they were, if you like "techno-political" decisions to make HTTP look like NNTP and mail [the Internet's existing standards for Usenet discussion groups and email]. I was trying to leverage as much of the existing technology and existing understanding out there as well. Same thing with HTML -- basing it on SGML, because that was the only common format people had talked about for hypertext at all.

continued
WN: You were essentially adding a separate layer, an additional application on top of this thing that was the

Internet, correct?

Berners-Lee: Well, the most important thing that was new was the idea of URI -- or URL [it was UDI back then, universal document identifier]. The idea that any piece of information anywhere should have an identifier, which will not only identify it, but allow you to get hold of it. That idea was the basic clue to the universality of the Web. That was the only thing I insisted upon.

WN: To make sure people understand about that new element, the URI. What was the URI, which became more commonly known as the URL?

Berners-Lee: [They are] these funny things that start http-colon-slash-slash, and then some gobbledygook which is the name of the document. The thing which is sometimes called the Web address, the thing which you find in shortened form painted on trucks and vegetables and all kinds of things now. Basically it identifies some piece of information out there on the Web.

WN: There was a short article in Wired magazine in 1993 about this thing it called "W3." Speaking of this newfangled idea of yours, the article read: "As soon as more and more client software becomes available, Berners-Lee expects more and more information to be woven into the Web." Did you have any sense at the time of the phenomenon you had touched off?

Berners-Lee: By '93, yes. It was never clear that it wouldn't just stop. Any time during that exponential growth, it could have stalled. I think we were never very confident until 1993.

WN: But the indications that you had done something that was catching fire were very clear by that point?

Berners-Lee: They were fairly clear, but even then [it could still have fragmented]. It could still fragment. There could still be a huge battle which leaves a big mess and [fragments the Web] into two pieces whenever a new

feature comes along. Everybody who runs a Web site knows we're not assured of compatibility, and we could end up with a split. For example, now the pressure is from the TV, PCs, PDAs. Different sized screens -- should they have different Webs? That was a very important initial assumption: Whatever the device you use for getting your information out, it should be the same information.

WN: In one of your recent interviews you mentioned that you considered, then abandoned, the idea of starting a dot-com company back in 1993. People may not appreciate the role you've played, steering the Web through the World Wide Web Consortium since then. What would have happened if the Web had not had a Tim Berners-Lee, or somebody in your stead, doing the work you've been doing to preserve its essence?

Berners-Lee: What could have happened is that you don't find simply a URL listed somewhere. You'd find a URL plus 'you must use this software' or 'you must get a particular piece of hardware in order to follow this link.' [If that happens], a URL is not enough. It's not enough to make a link to something. You have to say somewhere that you need [a] particular browser, you should probably be running on this sort of operating system, or this sort of hardware.

There was a huge amount of diversion that was happening just at that point -- '93 to '94. Every browser had its own flavor of HTML. So it was very difficult to know what you could put in a Web page and reliably have most of your readership see it. And that was a dire situation, which people at the Consortium worked on pretty hard.

continued
WN: So many ideals were put forth in 1993 and 1994, many of them by Wired magazine: There was the idea that the Web as an interactive medium would bring society together in ways no medium had before. Another idea was that ordinary

people could collaborate on new ideas. They could loosen the chains of one-to-many media, and bring about a new kind of electronic democracy.

In today's Internet IPO world, the Web seems to mean entrepreneurs slapping dot-com on the end of anything -- say, pet food -- hoping that venture capitalists will jump on board.

You say in your book that the Web is "not done," and it's the old ideals that dominate your vision for the future. Can we still get to a place where the Web can survive its current gold-rush mentality? What are the chances for your high-minded ideals to succeed?

Berners-Lee: It's not just high-minded ideals, but fun -- being able to play, doing the creativity in the Web, rather than doing it offline and then somehow compiling it into a Web page -- that sort of thing.

I think that it needs a lot more bits that I initially realized. I thought all we need is a decent, really intuitive editor for creating this stuff. But in fact, if you're going to use it collaboratively you have to very good access control too -- access control where you can create groups [of collaborating users].

And we didn't have all of that underneath. We didn't have the cryptography to actually make it secure, defined in standards. You can allow people to talk to each other, but then you have to guarantee that they know who's going to have access to the conversation and all that ... lots of other pieces which were necessary. That's one of the reasons why it didn't happen.

The question is interesting. It leads to the point that people are looking, at the moment, at applications which are built fairly thinly on the Web. The sort of thing where you can write the Perl script in an afternoon and produce a new Web site addressing a new market, or a new business model. And you can clean up, and you can go to the IPO remarkably quickly.

But nobody in that process has added to the ten-year, twenty-year vision of what the Web should fundamentally be and whether it should be changing. I hope that companies -- difficult for startups -- but I hope that large companies will continue to fund the research into the more distant future, and that the government will. And that we don't get this feeling that the Web's done. People keep asking me what I think of it now that it's done. Hence my protest: The Web is not done!

There are a lot of big challenging questions in there. There are a lot of things that are not easy to write in an afternoon. There are things that are going to take a long time.