

Distributed Environments for Large Data-Objects¹:

The Use of Public ATM Networks for Health Care Imaging Information Systems

*William Johnston, Jin Guojun, Gary Hoo, Case Larsen
Jason Lee, Brian Tierney, Mary Thompson*

Imaging and Distributed Computing Group
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, CA, 94720

Abstract

We are exploring the use of highly distributed architectures for all aspects of collecting, storing, analyzing, and otherwise manipulating and making generally available, large data-objects. These objects - typically the result of a single operational cycle of an instrument, and of sizes from tens of MBytes to tens of Gbytes - are the staple of modern analytical systems. It is further the case that many of the instrumentation systems that generate such data-objects are used by a diverse and geographically distributed community: Examples from the scientific community include physics and nuclear science high energy particle accelerators and detector systems, large electron microscopes, ultra-high brilliance X-ray sources, etc. There are correspondingly complex instrumentation systems in the health care community that generate large data-objects.

Apart from the basic issues of data collection, transport, and processing, our work is exploring the general issues of using the network to provide flexible high performance data storage strategies and location independent access to analysis systems.

In this paper, we describe some of the general aspects of such systems, and specifically describe the use of a shared (public) IP over ATM network to facilitate collection, storage, analysis, distribution, and delivery of several kinds of large data-objects.

1. This work is supported by the U. S. Dept. of Energy, Energy Research Division, Mathematical, Information, and Computational Sciences office (<http://www.er.doe.gov/production/octr/mics>), under contract DE-AC03-76SF00098 with the University of California, and by ARPA ITO (<http://ftp.arpa.mil/ResearchAreas.html>). Author's address: 50B-2239, Lawrence Berkeley National Laboratory, Berkeley, CA 94720. Tel: +1-510-486-5014, fax: +1-510-486-6363, wejohnston@lbl.gov, <http://www-itg.lbl.gov/~johnston>. This document is report LBNL-38849.

Contents

1.0 Introduction	3
2.0 Applications	4
2.1 The Health Care Information System Application	4
3.0 System Architecture	7
3.1 Data Collection	7
3.2 Network Storage: The Distributed-Parallel Storage System	9
3.2.1 DPSS Architecture	10
3.2.2 DPSS and the TerraVision Application	13
3.3 Data Management	15
3.3.1 The Health Care Application	16
3.4 Performance	17
3.4.1 Performance Issues Related to ATM	17
3.4.2 ATM Network to Workstation Performance	17
3.4.3 ATM Network Performance	19
3.5 Security Architectures for Large-Scale Remote Environments	21
3.5.1 Background	22
3.5.2 A Use-Condition Centered Security Model	23
3.5.3 Implementation of a Prototype Architecture	27
3.5.4 Relationship to Other Work	28
3.5.5 Prototype Application	29
4.0 Conclusions	29
5.0 Acknowledgments	29
6.0 Notes and References	30

1.0 Introduction

The advent of shared, widely available, high-speed networks is providing the potential for new approaches to the collection, storage, and analysis of large data-objects. Two examples of large data-object environments, that despite the very different application areas have much in common, are health care imaging information systems and atomic particle accelerator detector data systems.

Health care information, especially high-volume image data used for diagnostic purposes - e.g. X-ray CT, MRI, and cardio-angiography - are increasingly collected at tertiary (centralized) facilities, and may now be routinely stored and used at locations other than the point of collection. The importance of distributed storage is that a hospital (or any other instrumentation environment) may not be the best environment in which to maintain a large-scale digital storage system, and an affordable, easily accessible, high-bandwidth network can provide location independence for such storage. The importance of remote end-user access is that the health care professionals at the referring facility (frequently remote from the tertiary imaging facility) will have ready access to not only the image analyst's reports, but the original image data itself.

The nuclear science environment involves a detector at a particle accelerator. After several levels of data filtering, the first data of interest are "events" that are emitted at a rate of about one per second, and are about 20 MBy in size. Like many large-scale scientific experiments, the collaborators are located at laboratories and universities all over the country, as are the storage and computing resources needed to analyze the data.

This general scenario extends to other fields as well. In particular, the same basic infrastructure is required for remote access to large-scale scientific and analytical instruments, both for data handling and for direct, remote-user operation. See [Johnston95b].

In this paper we describe and illustrate a set of concepts that are contributing to a generalized, high performance, distributed information infrastructure, especially as it concerns the types of large data-objects generated in the scientific and medical environments. We will describe the general issues, architecture, and some system components that are currently in use to support distributed large data-objects. We describe two examples, one a health care information system that has been built, and is in prototype operation, and a nuclear science data handling architecture that is being designed and built.

The basic elements of the general architecture include:

- data collection and the instrument-network interface
- on-line storage that is distributed throughout the network (for both performance and reliability)
- processing elements - also distributed throughout the network - for various sorts of data analysis
- data management that provides for the automatic cataloguing (metadata generation) of the data being stored
- data access interfaces, including the application-data interface
- tertiary storage ("mass storage") management
- user access to all relevant aspects (application, data, metadata, data management)
- transparent security that provides access control for all of the systems components based on the resource-owner's policy

These elements all need to be provided with flexible, location-independent interfaces so that they can be freely moved around the network as required for operational or other logistical convenience.

2.0 Applications

2.1 The Health Care Information System Application

An example of a medical application that uses a distributed large data-object architecture is a system that provides for collection, storage, cataloguing, and playback of video-angiography images using a metropolitan area ATM network.

Cardio-angiography² is used to monitor and restore coronary blood flow, and though clinically effective, the required imaging systems and associated facilities are expensive. To minimize the cost of such procedures, health care providers are beginning to concentrate these services in a few high-volume tertiary care centers. Patients are typically referred to these centers by cardiologists operating at clinics or other hospitals; the centers then must communicate the results back to the local cardiologists as soon as possible after the procedure. The advantages of providing specialized services at distant tertiary centers are significantly reduced if the medical information obtained during the procedure is not delivered rapidly and accurately to the referring physician at the patient's home facility. The delivery systems currently used to transfer patient information between facilities include interoffice mail, U.S. Mail, fax machine, telephone, and courier. Often these systems are inadequate and potentially could introduce delays in patient care. (See [Kaiser-CalREN].)

Using a shared, metropolitan area, ATM network, and a high-speed distributed data handling system, video sequences and still image are collected from the video-angiography imaging systems, stored, and accessed by a remote user. The image data are sent through the network to storage and analysis systems, as well as directly to the users at clinic sites. Thus, data can be stored and catalogued for later use, data can be delivered live from the imaging device to remote clinics in real-time, or these data flows can all be done simultaneously. Whether the storage servers are local or distributed around the network is entirely a function of the optimal logistics: There are arguments in regional health care information systems for centralized storage facilities away from the hospital environment, even though the architecture is that of a distributed system. (See [Johnston93].)

Figure 1 illustrates the configuration of the health care application. This application is a joint project of Lawrence Berkeley National Laboratory, Kaiser Permanente, Philips Research (Palo Alto), and the Pacific Bell CalREN program. (See [Kaiser-CalREN].)

The network aspects of this health care project are closely related to the BAGNet CalREN project (see [BAGNet-1] and [Wiltzius]). BAGNet was the first, and largest, of the CalREN ATM projects, and it established the basic TCP/IP over ATM infrastructure used in this and several other CalREN projects. The physical infrastructure is illustrated in Figure 2, and the logical and network infrastructure are described in [Johnston95c].

2. Cardio-angiography imaging involves a two plane, X-ray video imaging system that produces from several to tens of minutes of digital video sequences for each patient study for each patient session. The digital video is organized as tens of data-objects, each of which are of the order of 100 MBytes.

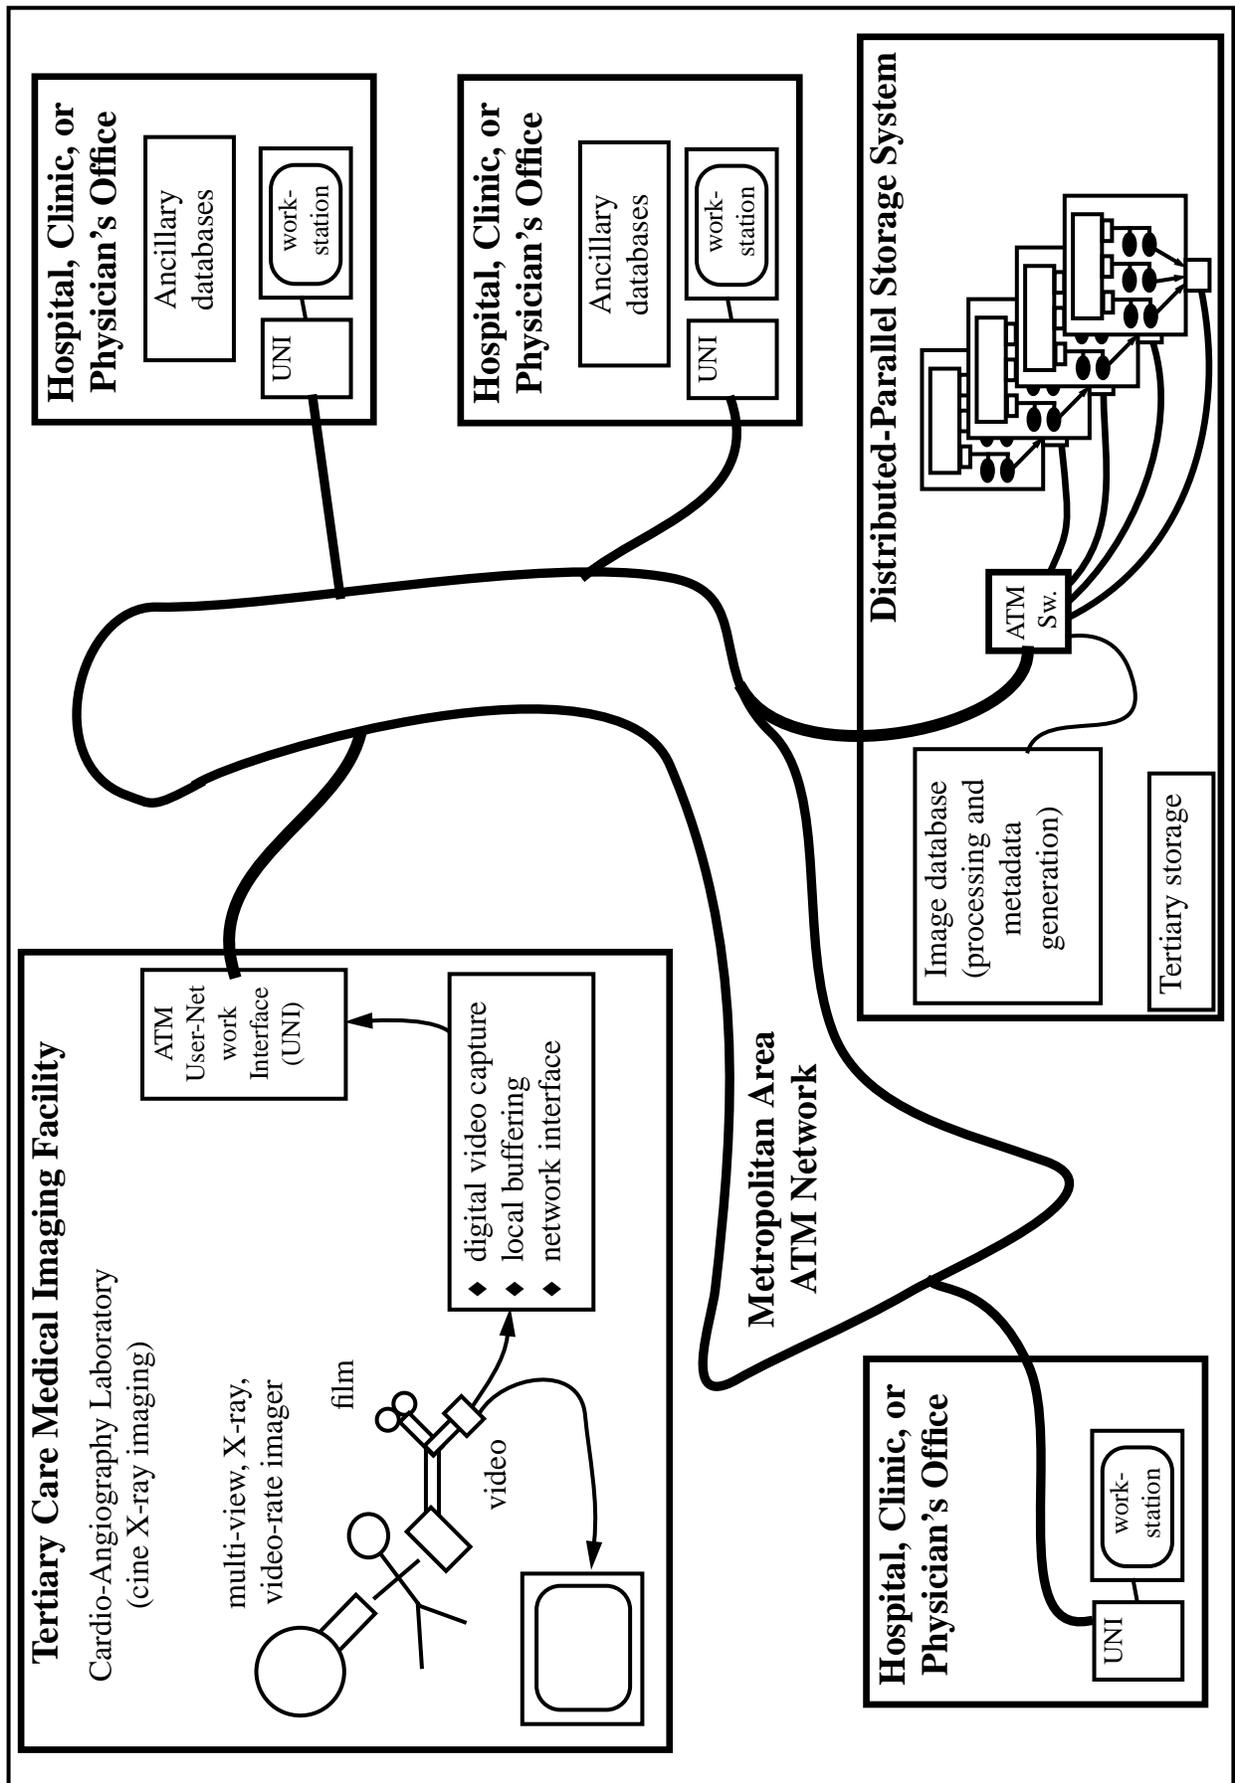


Figure 1 Distributed Large Data-Object Medical Imaging Application

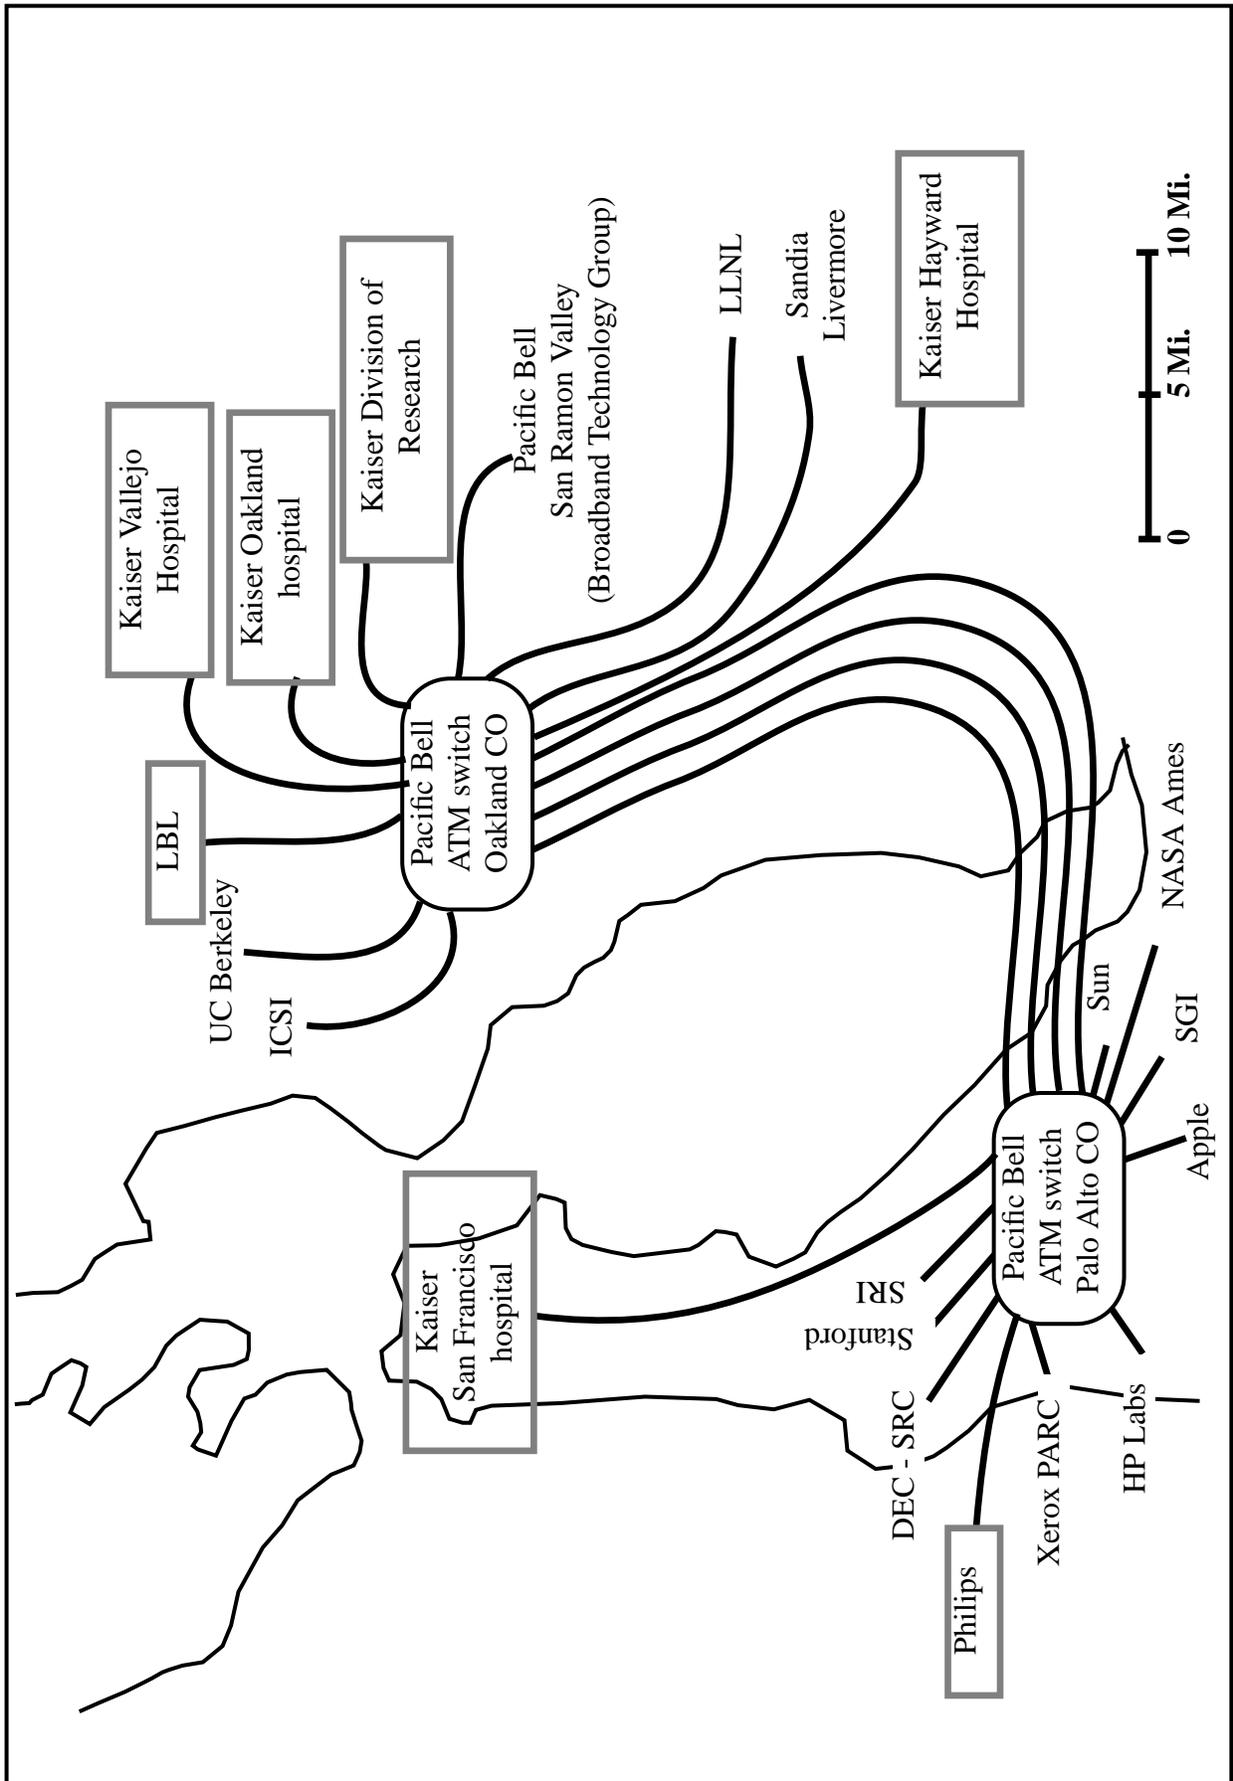


Figure 2 The Pacific Bell, S. F. Bay Area, Metropolitan Area ATM Network: BAGNet and Kaiser CalREN Sites

Figure 3 illustrates the health care application data flow, and once again, though we are describing a specific application, this data flow architecture is quite general, and we are using an almost identical architecture in several other projects involving the collection, cataloguing, storage, and access of data from several types of (unrelated) scientific experiment scenarios.

The angiography data is being collected directly from a Philips scanner by a computer system in the San Francisco Kaiser Cardiac Catheterization Laboratory. This computing system, in turn, is attached to a Pacific Bell, OC-3, metropolitan area ATM network. When the data collection for a patient is complete (about once every 20-40 minutes), about 500 MBy of digital video data is sent across the network to LBNL (in Berkeley) and stored first on an experimental, high-speed, network-based multimedia storage system (the DPSS) and then is archived to a mass storage system. This process goes on 8-10 hours a day.

The data is catalogued in a World Wide Web based image database system (“ImgLib”). Use of the data involves access through ImgLib to locate the data sets of interest, and then to invoke a viewing application.

Department-level Web-based patient databases can refer directly to the data in ImgLib without duplicating the data, or being concerned about tertiary storage management (which is handled by ImgLib).

3.0 System Architecture

3.1 Data Collection

Instrumentation systems are at the front-end of many distributed large data-object environments. Examples include particle accelerator detectors, Earth environment monitoring satellites, and medical imaging systems. These sources generate essentially continuous data streams, but ones that have “natural” boundaries that define “objects”: Video sources (like the cardio-angiography studies where data is collected in bursts); large single images; detector “event” records etc. We make this “natural boundary” characterization to distinguish these data types from continuous-media data (such as video and audio multimedia streams): Even though the average data rates may be comparable, or even higher in the large data-object environment, the temporal synchronization requirements and steady-state nature of multimedia streams give them a somewhat different character.

For many of the instrumentation systems that we are interested in, one of the primary issues is getting the data out of the instruments and on to the network. One of the circumstances that has led to both the interest in, and practicality of, the architecture being described here is that general purpose workstations now have memory and I/O bus structures that are fast enough to acquire, structure, and send to the network, significant bandwidth data streams. This is an important capability because it means that the only special hardware that is required to bring instruments on-line is the interface between the workstation and the data source, and even this interface may be able to be done in “software” using off-the-shelf DSP-based I/O boards.

The frontend workstation acquires the raw data, formats it as “objects” by adding or using metadata from the experiment environment, and then sends the objects into the distributed environment. The data collection workstation frequently also serves as a buffer so that brief interpretations or slow-downs in the network do not result in loss of data. This frontend architecture is illustrated in Figure 4.

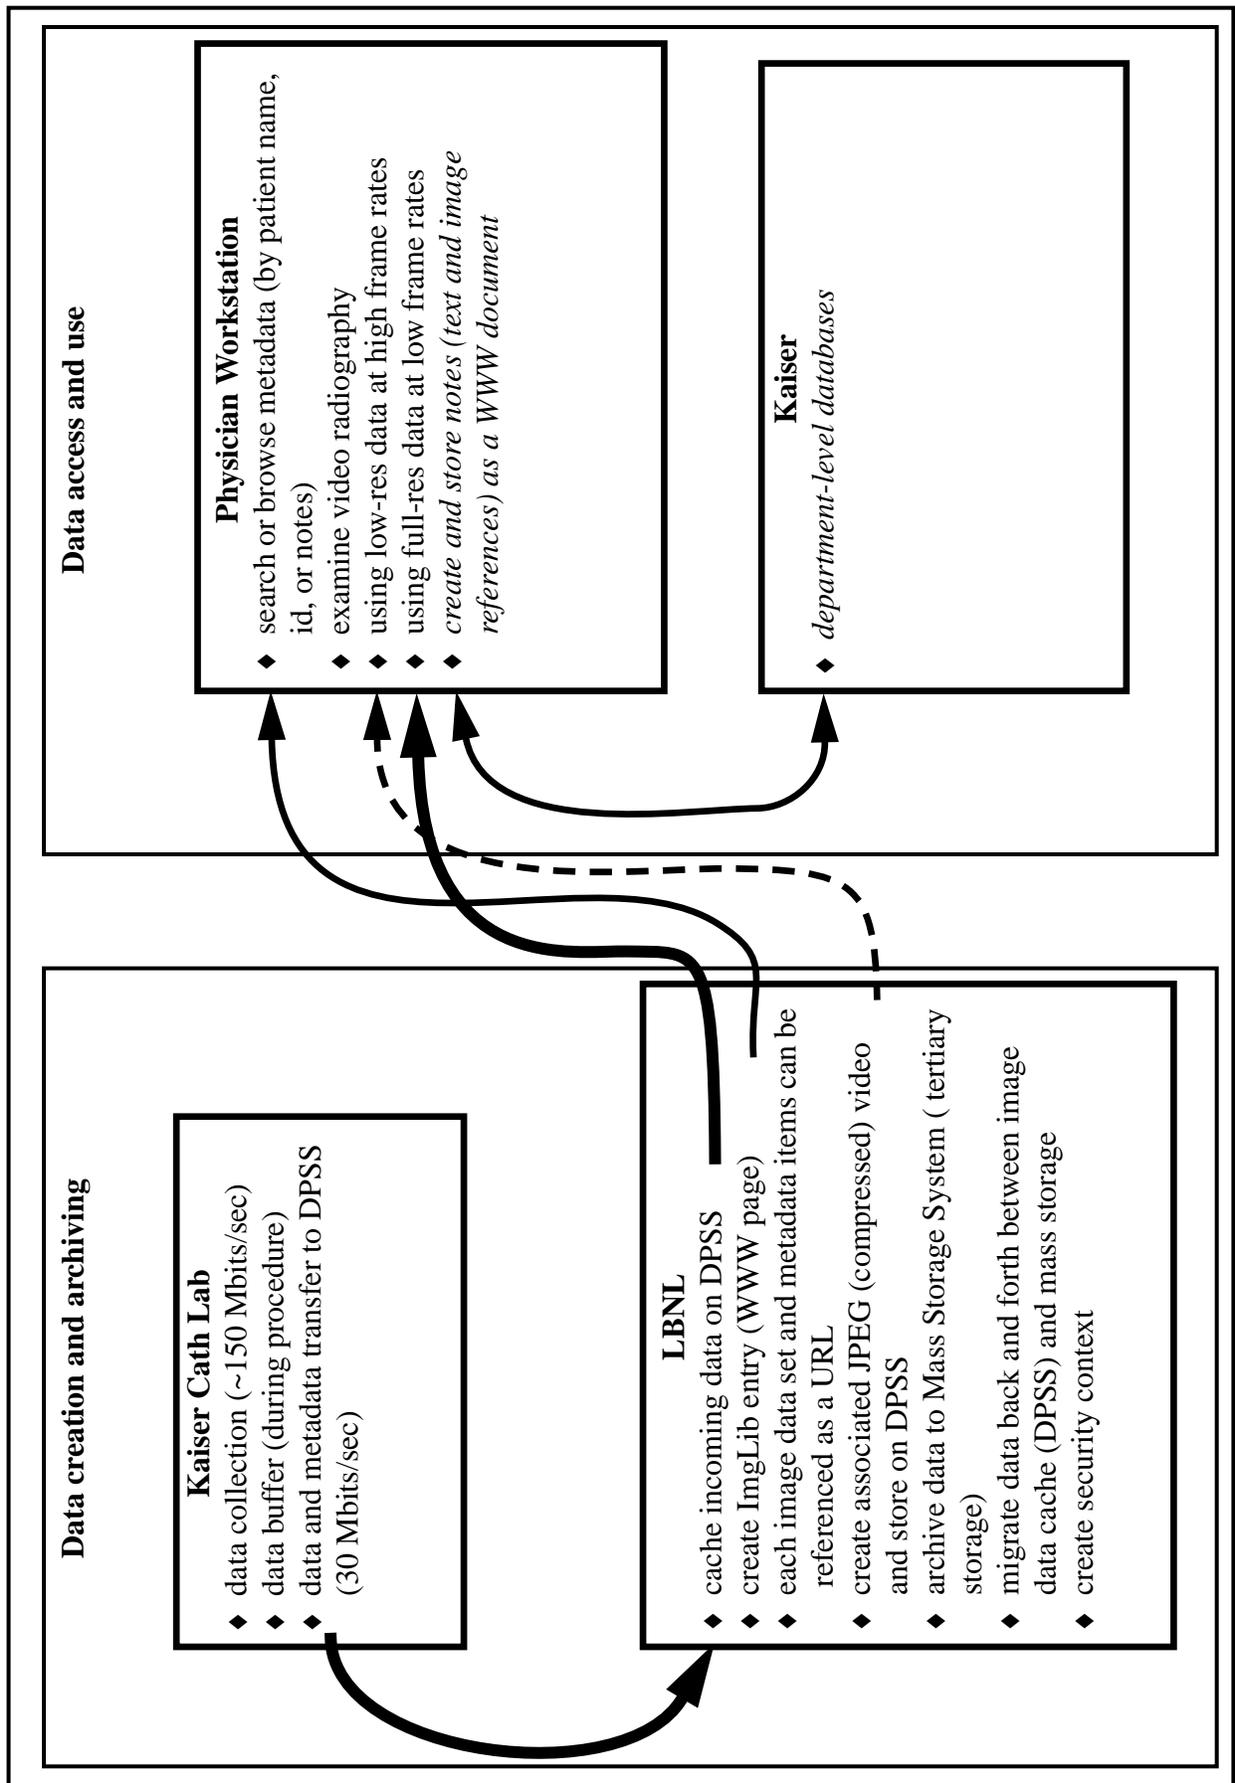


Figure 3 Kaiser Project Data Flow (A typical data flow for distributed large data-object systems)

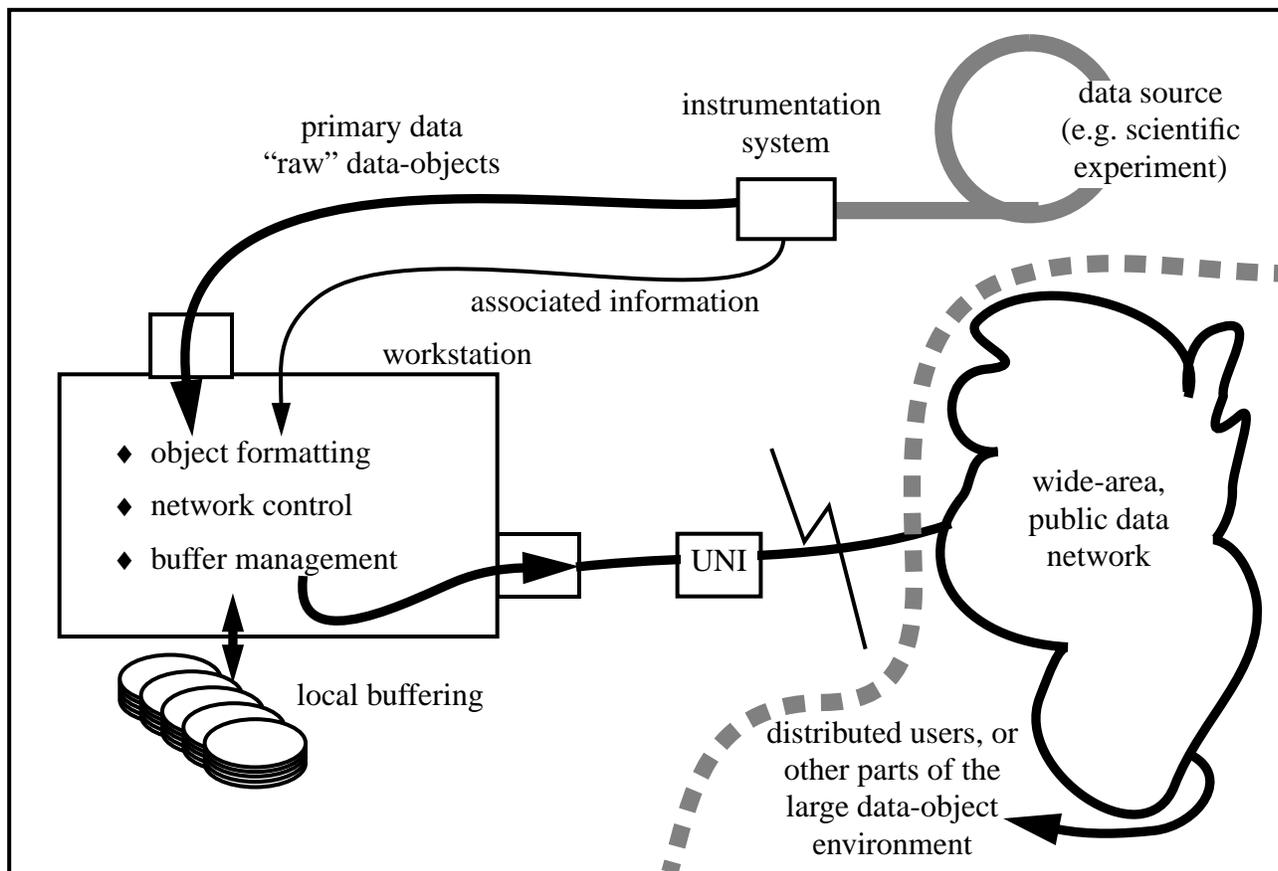


Figure 4 Architecture of the Frontend of a Distributed, Large Data-Object Environment

3.2 Network Storage: The Distributed-Parallel Storage System

Widely distributed network storage systems are an essential component of a network-based large data-object environment.

Distributing the components of a storage system throughout the network increases its capacity, reliability, performance, and security. Usable capacity increases in conjunction with a widely deployed, generalized security infrastructure that can support dynamic construction of systems through brokering and automated acquisition of resources. (See [Johnston96a].) Reliability increases because storage systems that can be configured from components that have as little as possible in common (e.g. location) provide the resilience that comes from independence. Performance is increased by the combined characteristics of parallel operation of many sub-components, and the independent data paths provided by a large network infrastructure. Security is also potentially increased by having many independent components each of which has local and independent enforcement mechanisms that can limit the scope of a security breach.

The Distributed-Parallel Storage System (DPSS, aka ISS) is an experimental system in which we are developing, implementing, and testing these ideas. In most configurations, it is used as a network-striped disk array designed to supply and consume high-speed data streams to and from other processes in the network. (See [DPSS].)

3.2.1 DPSS Architecture

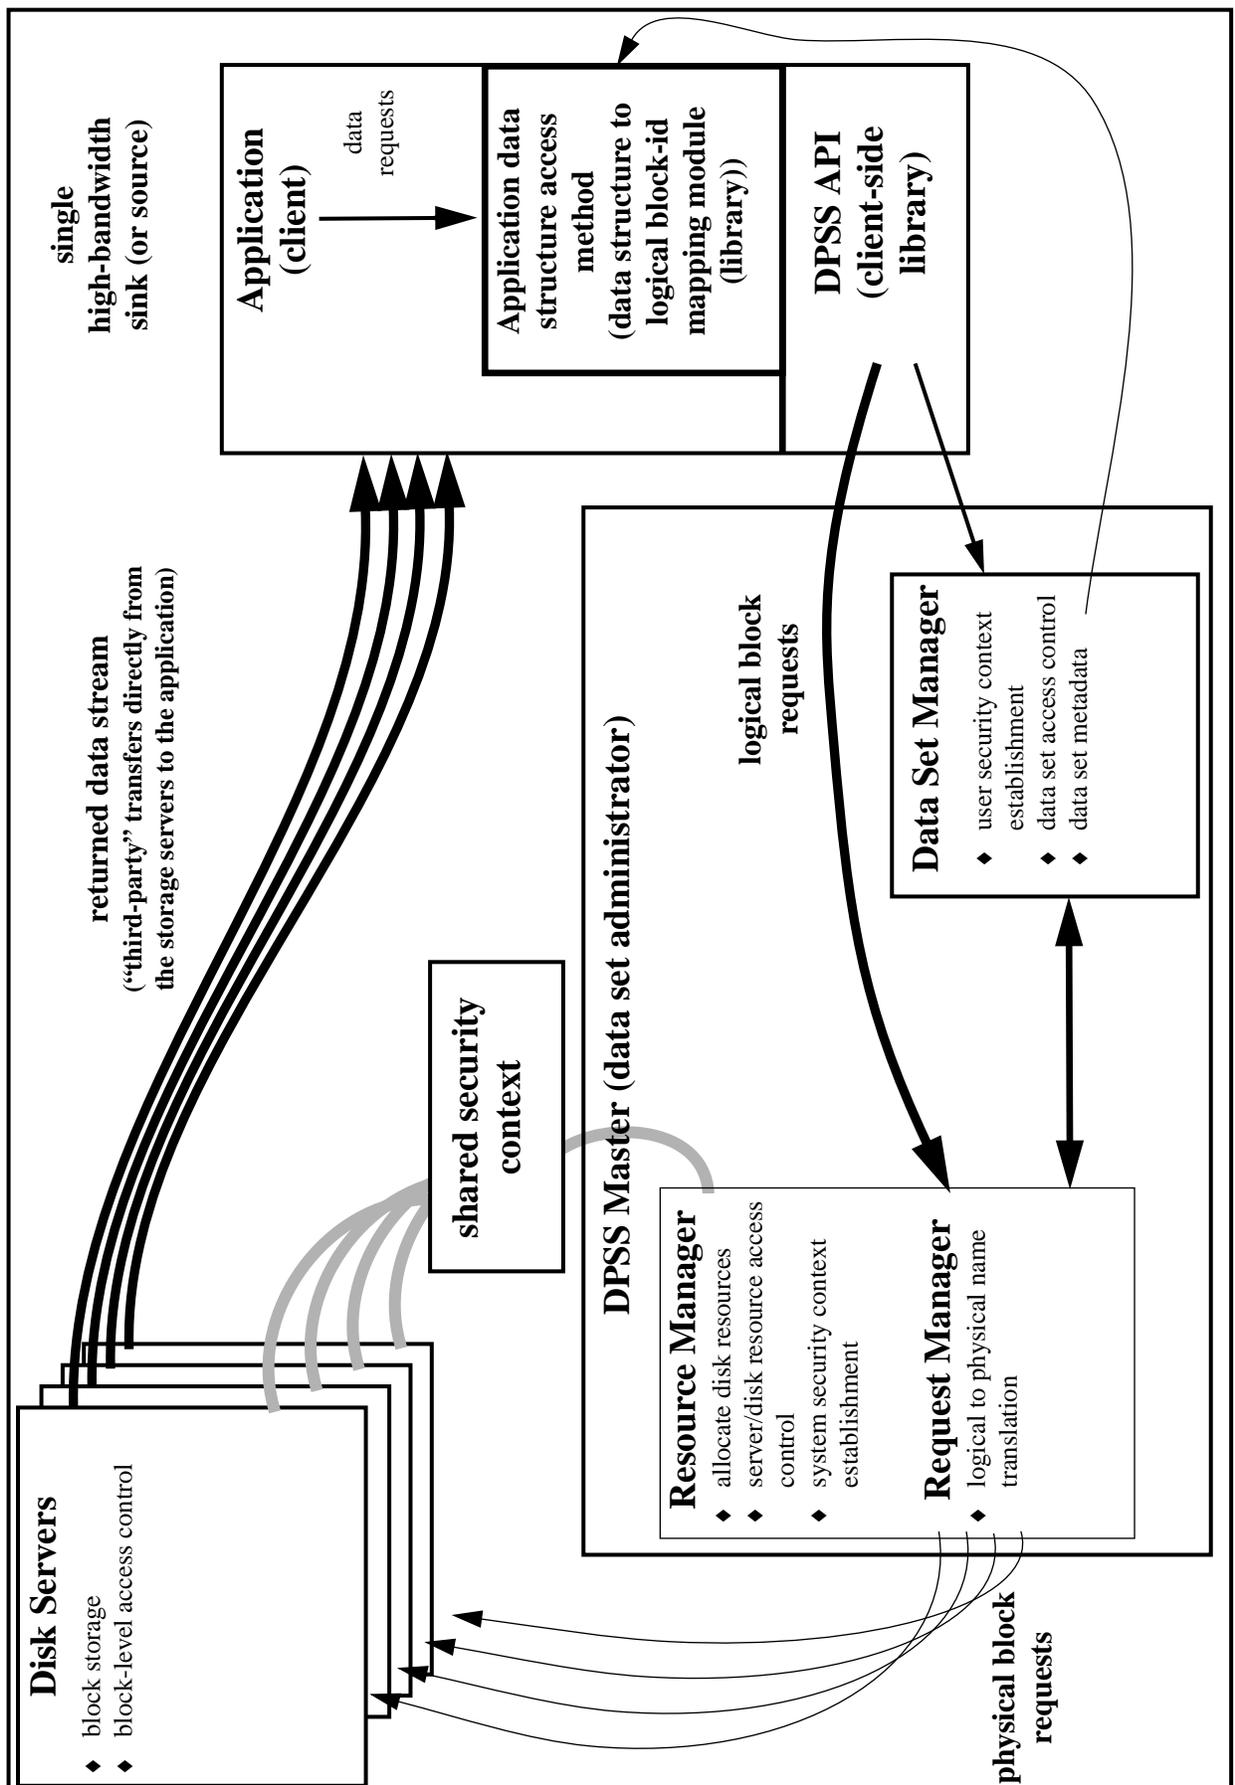
The DPSS is essentially a “logical block” server whose functional components are distributed across a wide-area network. (See Figure 5 illustrating the DPSS architecture.) The DPSS uses parallel operation of distributed servers to supply image streams fast enough to enable various multi-user, “real-time”, virtual reality-like applications in an Internet / ATM environment. There is no inherent organization to the blocks, and in particular, they would never be organized sequentially on a server. The data organization is determined by the application as a function of data type and access patterns, and is implemented during the data load process. The usual goal of the data organization is that data is declustered (dispersed in such a way that as many system elements as possible can operate simultaneously to satisfy a given request) across both disks and servers. This strategy allows a large collection of disks to seek in parallel, and all servers to send the resulting data to the application in parallel, enabling the DPSS to perform as a high-speed data server.

The implementation is based on the use of multiple low-cost, medium-speed disk servers which use the network to aggregate multiple server outputs for high performance applications. To achieve high performance all types of parallelism are exploited, including those available at the level of the disks, controllers, processors / memory banks, servers, and the network (See Figure 6 illustrating the DPSS implementation).

At the application level, the DPSS is a semi-persistent cache of named objects, and at the storage level it is a logical block server. Although not strictly part of the DPSS architecture, the system is usually used with an application agent called a “data set structure method”. This component provides an object-like encapsulation of complex user-level data structures so that the application does not have to retain this information for each different data set. The function and interface of the structure methods are left to the application domain, but one simple example is for video data. In this case the structure method allows applications data structure to request data by “frame” number and “type” (e.g. in the video-angiography application, there are two images - 90 degree offset views - associated with every frame). The structure server converts the application requests into logical block requests. These logical block requests are then sent to the DPSS “data set administrator” (DSA). The DSA maintains data set definitions, and is responsible for mapping the logical block requests to physical block requests. The DPSS DSA also deals with interactions with the storage servers to determine available storage (a storage server is an independent entity and may deal with several DPSS administrator) and to establish the “security context” that provides the scope of control for various resources.

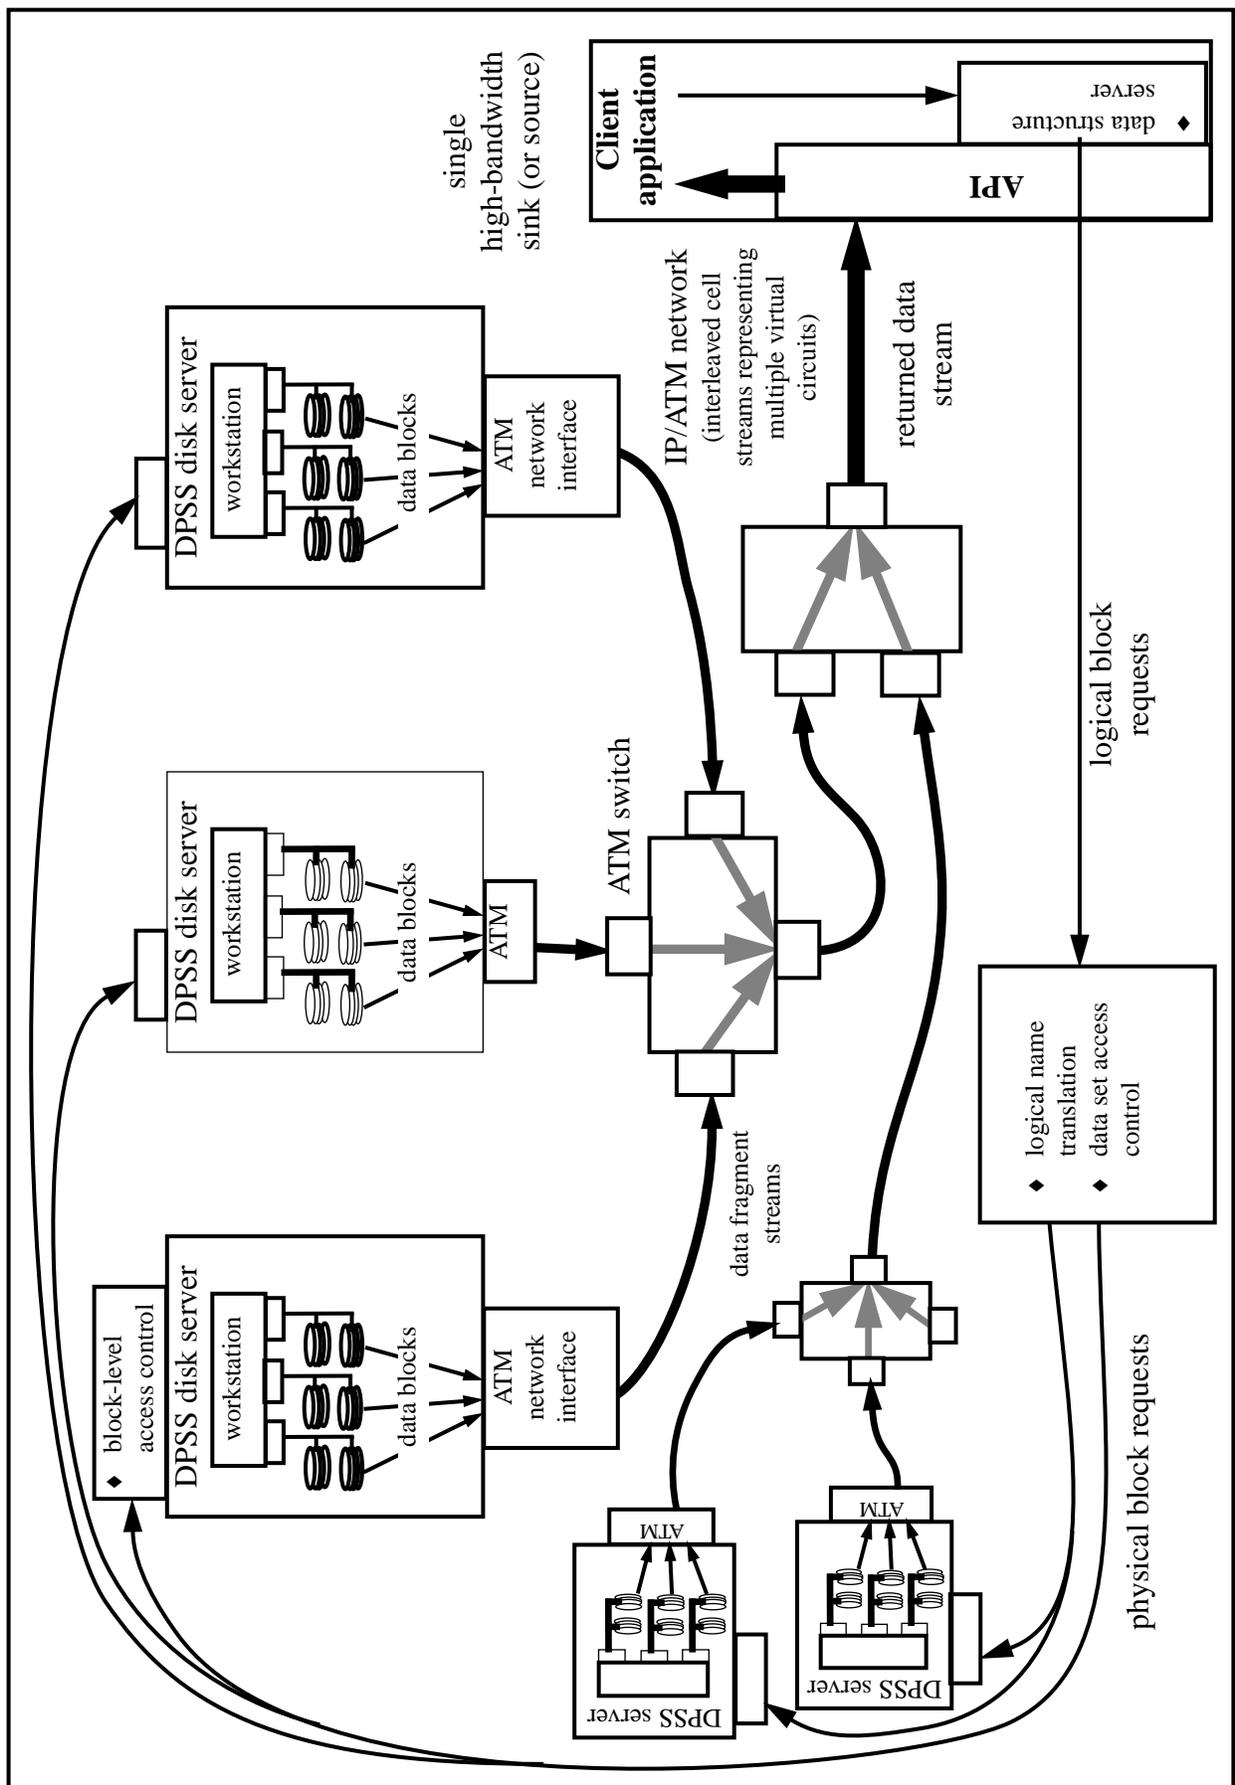
The security model for the DPSS involves accommodating several different resource owners. The context established between the DSA and the disk/storage servers reflects an agreement between the owner of the physical resources (disks) and an agent that is providing storage to a user community. This context enforces the disk usage agreement. The separate context established between the DSA and the users reflects the use-conditions imposed by the data “owner”, and provides for ensuring access control that enforces those use-conditions.

The overall data flow involves “third-party” transfers from the storage servers directly to the data-consuming application (a model used by most high performance storage systems). Thus, the application requests data, these requests are eventually translated to physical block addresses (server name, disk number, and disk block), and the servers deliver data directly to the application. (The servers and the application host are made known to each other during the data set “open” process, and sufficient information from the original request is carried along so that the application can identify the data provided by the servers.)



Distributed-Parallel Storage System Architecture

Figure 5



Distributed-Parallel Storage System Implementation

Figure 6

The DPSS is not an inherently reliable tertiary storage system, though the potential exists for this capability by adding tertiary backing storage on the storage servers.

3.2.2 DPSS and the TerraVision Application

Figure 7 illustrates the configuration and operation of the DPSS in the ARPA-funded MAGIC high-speed ATM testbed [MAGIC] where the DPSS is being developed. The primary characteristics of this environment are the various networks (MAGIC backbone, Spartan [Spartan], and NTON [NTON]), the TerraVision terrain visualization application [TerraVision], and the model of a data curator (EDC). All components of the DPSS are geographically dispersed in MAGIC via the wide-area, high-speed, network infrastructure.

3.3 Data Management

In any scenario where data is generated in large volumes and with high throughput, and especially in a distributed environment where the people generating the data and the people cataloguing and using the data are geographically separated, a key issue is the automatic generation of at least minimal metadata, and the cataloguing of the data and the metadata as the data is received (or as close to real time as possible).

Such a capability is needed to help researchers organize, browse and search through data collections, especially collections where the original data are stored off-line in a mass storage system. The system should also facilitate co-operative research by allowing access to the data by specified users at local and remote sites, both by enabling immediate examination of the data, and thought incorporation of the data into other databases or documents.

Our approach to this capability (called “ImgLib”) is to use the World Wide Web tools to provide a semi-automatic cataloguing of incoming data. This is done by extracting associated metadata and converting it into text records, generating auxiliary metadata and derived data, and combining these into Web documents. (See [ImgLib].)

In the case of image-like data, an ImgLib catalogue “entry” consists of (derived) thumbnail images and textual indexing material, together with pointers to the original images, video, or other data-objects (whose permanent storage is typically off-line). The indexing material is kept on-line on a system that is running the ImgLib Web Server. The ImgLib server also directly manages a DPSS-based on-line cache, and manages migration of data from tertiary storage to cache.

A typical ImgLib usage scenario involves creating a top level collection (a branch of a directory tree of html documents) on the server machine. Once the initial organization is established, the collection owner is able to manage the collection from his local workstation via the Web “forms” interfaces. ImgLib provides the following functions:

- Creates (from the original data-objects) and stores a set of “browsable” (e.g. “thumbnail”) images that can be easily viewed;
- Creates a text description file associated with each data-object and indexes over these files;
- Allows the collection to be laid out in hierarchical sub-collections which can be independent of how the original data-objects are organized or stored;
- Provides automatic and transparent management of the tertiary storage and the migration of data-objects from tertiary storage to on-line cache;
- Provides mechanisms to build, on-the-fly, subsets of the collection of data-objects based on keyword or metadata field searches or sub-collections defined by branches of the organizational hierarchy;
- Provides browsing tools to present a “view” of the subsets represented by thumbnail images and the associated textual metadata.

Figure 8 illustrates some of these points in the context of a health care video imaging application. It shows the results of a search on the textual metadata. The information about the data-objects that result from the search is shown as an associated collection of thumbnails, associated pointers to several other types of derived data (including a “movie” representation), and a pointer to the original data-object. The data-objects are kept on tertiary storage (a tape-robot based mass storage system in this case), and there is an option for forcing migration of data back to the on-line cache if the data of

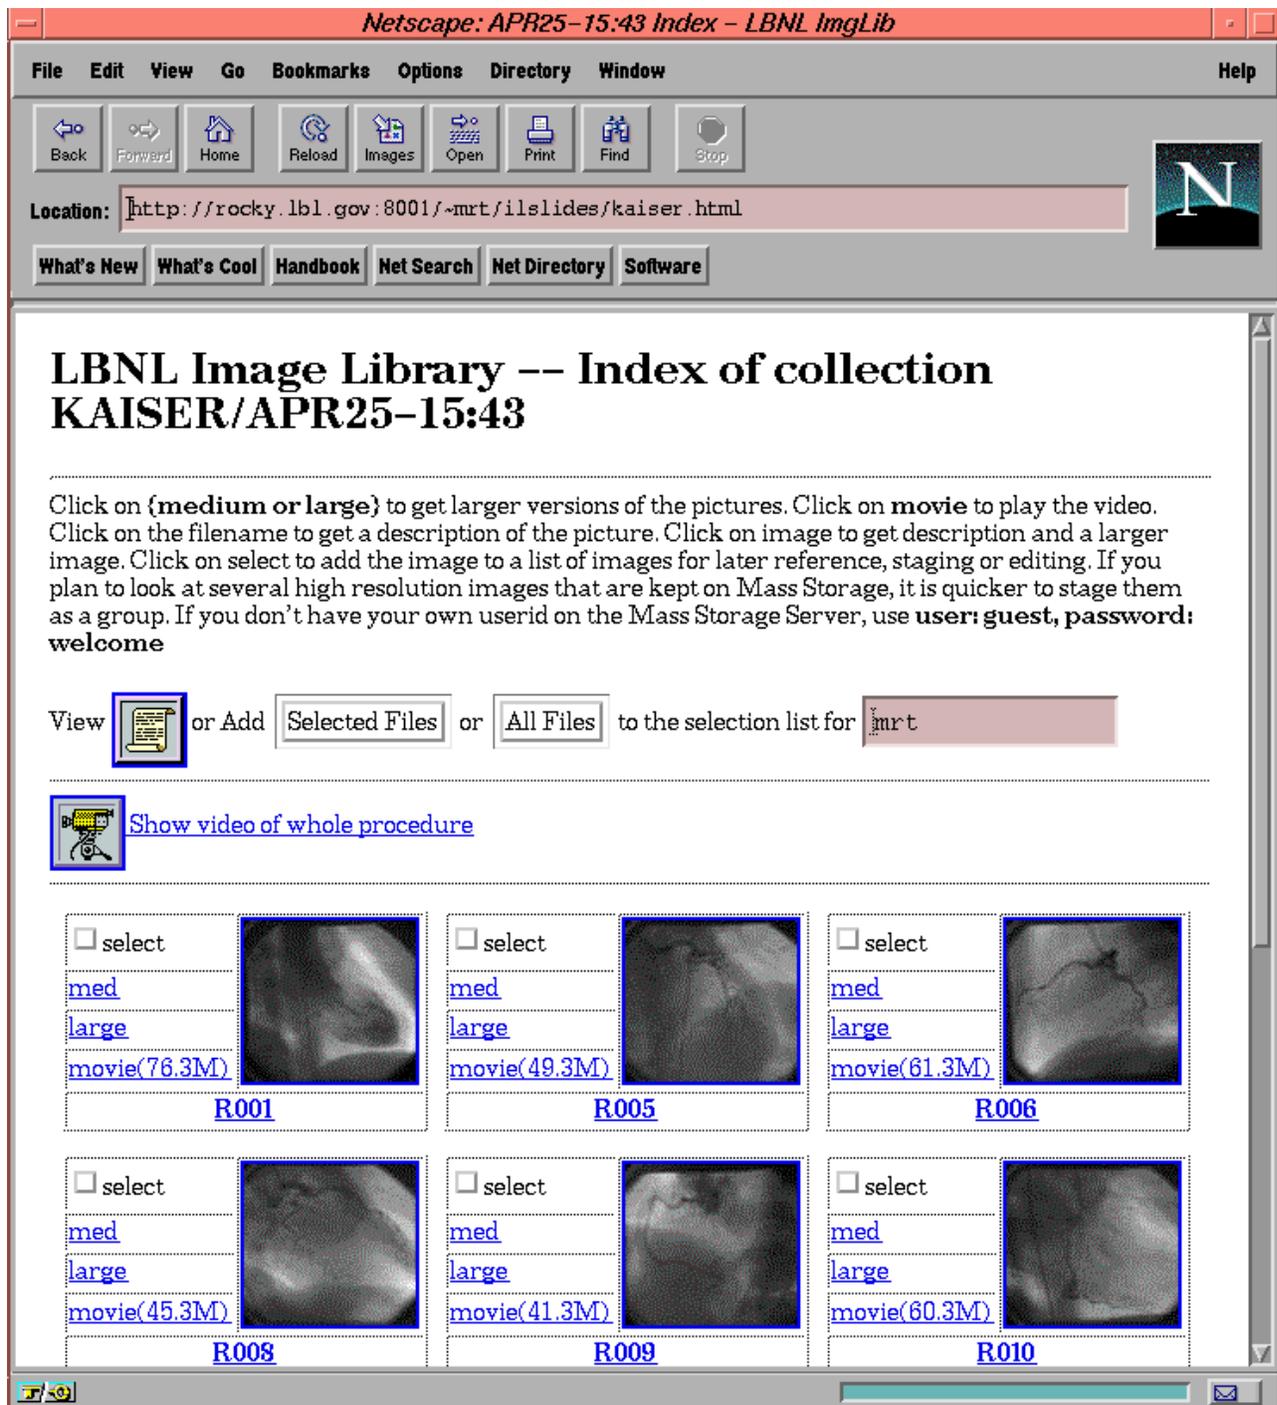


Figure 8 Large Data-Object Management for Video-like Objects

interest is not already there. In this example, a JPEG “movie” is the derived representation. (The original video data cannot be compressed and requires a special application to view.)

3.3.1 The Health Care Application

In the case of the cardio-angiography application, the video data is processed and then indexed using the method described above. Subsequently, physicians at Kaiser Oakland Hospital, Kaiser Division of

Research, and Kaiser Vallejo Hospital can locate images and video sequences through the Web interface, and study those sequences from their workstations. The volume of original digital video data is sufficiently large that it is not practical to sent it across the network for viewing at full frame rates on typical physician workstations. Therefore the video data is viewed with a special video player that first uses a compressed version of the data to play the video-angiography sequences at video rates, and then to display a subset of the original data. The compressed sequences are used as a “preview” to locate the individual angiography images of interest. As the video player is “slowed down” by the physician (for more detailed examination of the individual frames) and the frame rate reaches three or four frames per second, the player switches to displaying the uncompressed digital video directly from the DPSS network storage system. (See Figure 1.) This approach uses the network first to transport the compressed video sequence, and then, in real-time, only those full-resolution frames needed for close examination. This avoids having to move the very large amount of un-compressed digital video around the network many times and makes it possible for physicians to call up data “on the spur of the moment” and examine it on commonly available workstations that are not able to locally store the tens of Gigabytes of data needed for a full angiography study.

3.4 Performance

Network performance and characteristics are obviously an important factor in the architecture that we are describing. We have developed several performance monitoring strategies for characterizing network performance when they are used for high data rate applications.

3.4.1 Performance Issues Related to ATM

There are virtually no behavioral aspects of an ATM “network” that can be taken for granted, even in an end-to-end ATM network. By “network” we mean the end-to-end data path from the transport API through the host network protocol (TCP/IP) software, the host network adaptors and their device drivers, the many different kinds of ATM switches and physical link bandwidths, and then up through the corresponding software stack on the receiver. Further, the behavior of different elements at similar places in the network architecture can be quite different because they are implemented in different ways. The combination of these aspects can lead to complex and unpredictable network behavior.

The next two sections provide some information and observations on this issue. We only address issues in the flow of data after the host-to-host connection is established. (That is, we do not address routing and signaling.)

3.4.2 ATM Network to Workstation Performance

The architecture and resulting environment that we are describing is a “data movement intensive” environment. That is, in the normal course of dealing with large data-objects in a fully distributed environment, large amounts of data are constantly on the move. This circumstance generates many the R&D issues from workstation network performance to security, and we discuss a few of them here.

In a scenario where parallelism is used to provide high-performance, then one of the important attributes is the independence of the parallel operating elements. One of the reasons that the DPSS has been so successful as a high performance storage system is that using wide-area networks increases the independence of the servers (until the final point of delivery).

One of the problems that we have encountered is that independence within the workstation platforms is much more difficult to achieve. Fortunately, platform hardware and software designers have put considerable effort into making key sub-systems operate independently. So, for example, we have found that the parallel operation of disks on controllers, controllers on busses, and memory sub-systems all do a reasonably good job of increasing performance through parallel operation (though they may still be slow in absolute terms).

On the other hand, this is not true for network interfaces. That is, through progress is being made, multiple network interfaces have, so far, generally done a poor job of providing increased throughput by parallel operation. For ATM interfaces, this is due in part to the newness of the technology (and corresponding lack of experience in designing the supporting system hardware and software), and in part due to the fact that prior to the commercial interest in video file servers, striping across network interfaces was not a common practice.

The following synopsis describes the progress in this area (which is very important in high-speed distributed applications). Some of the progress is due to the maturing of ATM network code stacks, and some is due to several years of interaction with the workstation vendors pointing out the problems.

The evolution of the IP over ATM infrastructure:

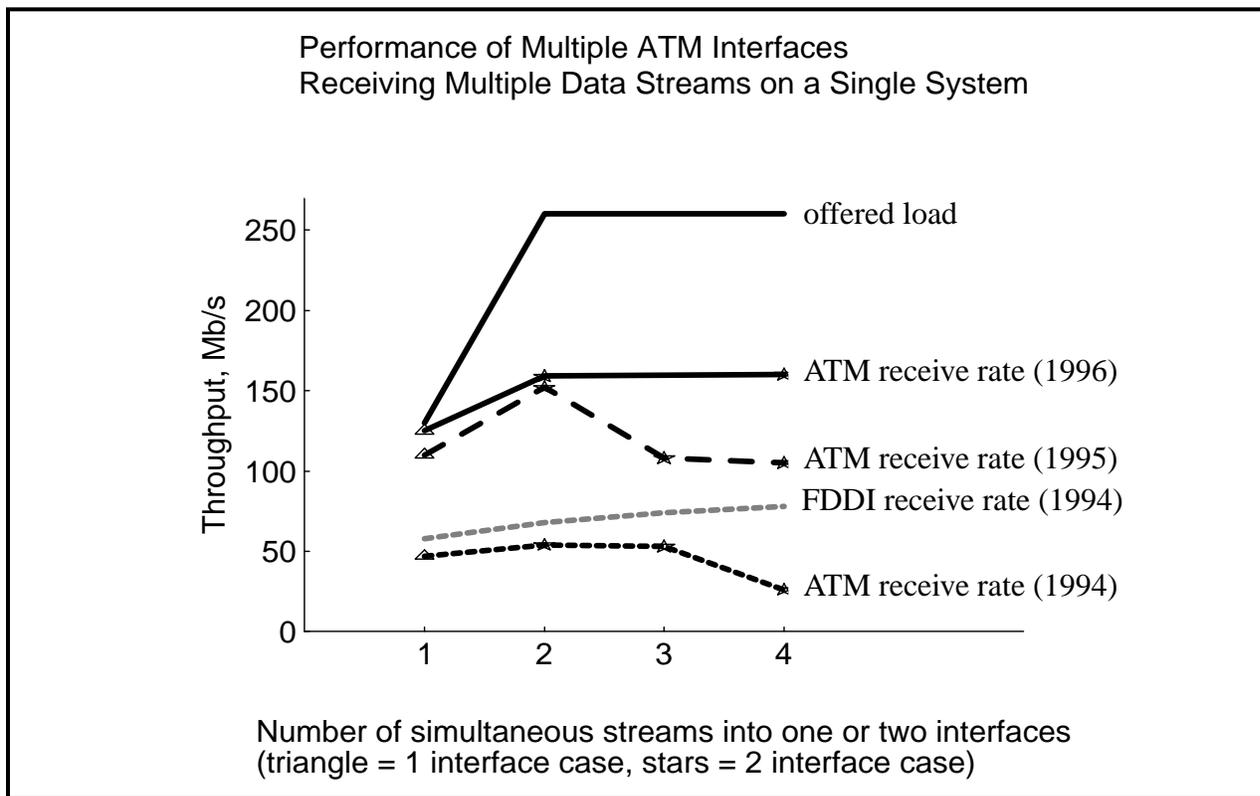


Figure 9 Evolution of the Performance of Workstation ATM Network Interfaces

- 1994
 - single interfaces were slow
 - multiple interfaces were no faster (poor independence of system data paths - nothing was multi-threaded)

- switches dropped cells silently
- 1995
 - single interfaces become faster
 - multiple interface data path independence improved
 - switches still dropped cells silently
- 1996
 - single interfaces and systems are approaching wire speed
 - multiple interfaces better still, but not ideal
 - switches now report cell loss

3.4.3 ATM Network Performance

3.4.3.1 Uncongested Behavior

Application Requirements:

To provide a specific example, we consider the Kaiser CalREN project and environment ([Kaiser-CalREN]).

Medical imaging applications are a “systems” problem - no one requirement sets all of the necessary parameters). Distributed storage, processing, and distribution of video-angiography large data-objects provides an example of an application constrained by the “fair share” use of a public data network. Data is collected and stored locally during a short recording period, after which it is transmitted to a network-based storage system for subsequent processing, retrieval, and review at other locations. The quantitative characteristics are:

- video-angiography \equiv x-ray video camera in a bi-plane (two view) configuration (Phillips DCI Scanner)
- data is collected at $30 \text{ fr/sec} \times (512 \times 512 \text{ pixels} \times 8 \text{ bits/pixel} \times 2 \text{ views})$ (monochrome video, no compression) = 126 Mbits/sec (one of several relevant parameters)

Four GBytes of data is collected in a local buffer in about four minutes in the form of data-objects that range in size from 50-300 MBytes. Following the data collection, there are about twenty minutes in which to off-load local storage before the next collection period.

Four GBytes in 20 minutes is a rate of about 25Mbits/sec of user-level data on the network, or, accounting for the various protocol overheads, approximately $25 \times 1.4 \approx 34 \text{ Mbits/sec}$ at the ATM level.

34 Mbits/sec is about 23% of an OC-3 network, and it was our initial estimate of the maximum reasonable use of a shared network, in this case Pacific Bell’s SF Bay Area ATM network - an ATM OC-3 network which has 30-40 attached sites.

Observations:

We ran a series of throughput tests to determine suitability of the Pacific Bell Metropolitan Area ATM for video-angiography transmission.

The experiments consisted of sending 25 Mbits/sec, with a duty cycle of 40 minutes on, 20 minutes off, 24 hours a day, and were conducted over several weeks.

Although this network was basically uncongested, several of the experiments were timed to correspond to some high resolution video and audio multicast sessions that generated 5-10 Mbits/sec of “continuous media” (steady state) traffic in the network, in addition to whatever other activity might have been occurring in the network.

The experiments were done by putting the source host ATM adapter into “continuous bit rate” mode, but sending data via IP over AAL-5. This ensured that we transmitted a “true” 25 Mbits/sec, and did not inject high-bandwidth “bursts” of data into the network (as would have been the case if we had rate-limited the TCP writes to an “average” of 25 Mbits/sec).

As can be seen in Figure 10, to the level of about 1%, there was no variation in our ability to send 25 Mbits/sec of data for any of the several-week-long periods during which the throughput experiments were run. Figure 10 shows typical results that include some times with high rate multicast traffic was present (in particular, Friday afternoon).

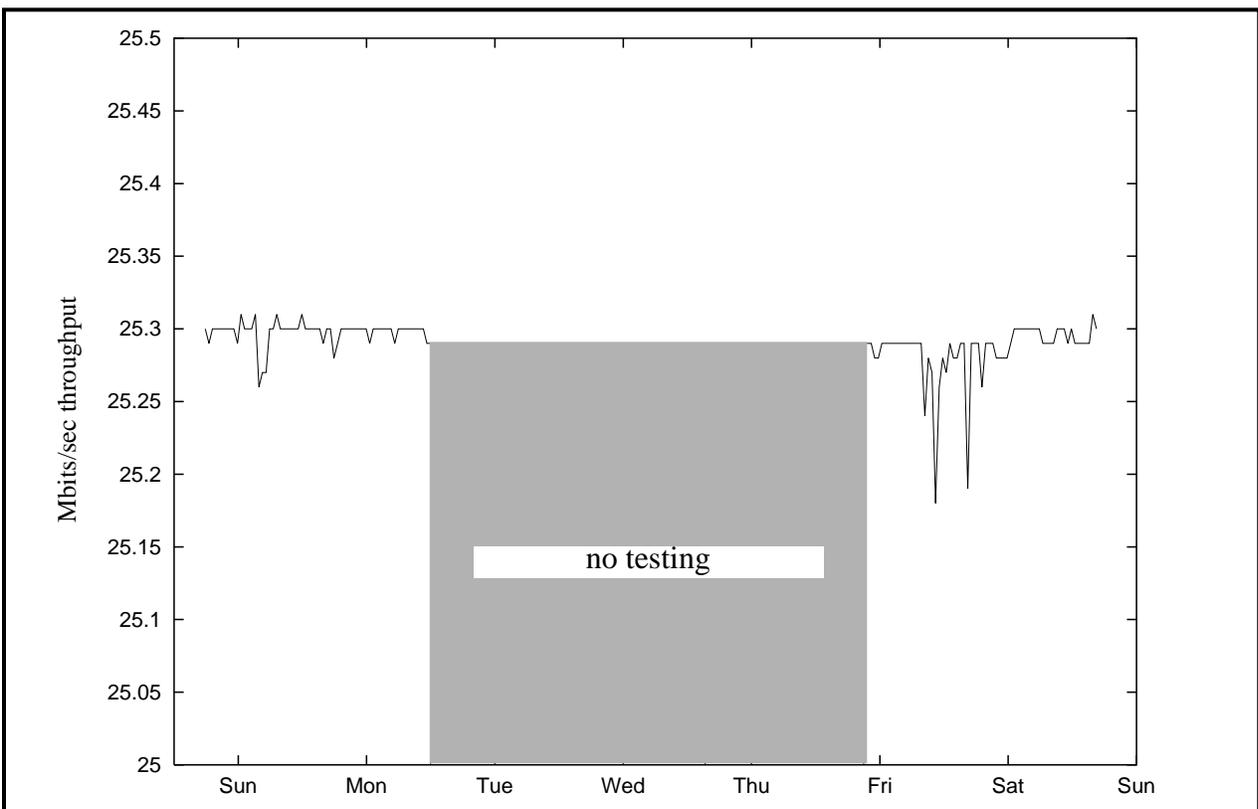


Figure 10 Throughput Testing to Determine the Suitability of the Pacific Bell, SF Metropolitan Area ATM for Video Angiography Transmission

3.4.3.2 Congested Behavior

The impact of congested ATM networks on TCP transport (which is used for all of the data transfers described here) can be very complex. Today’s ATM networks and components are still relatively immature and with every new generation of switch software and hardware the congested behavior can

change. The most severe impacts occur in circumstances when the characteristics of the various ATM links and network components conspire to defeat TCP's otherwise well-tuned congestion control mechanisms. TCP has a suite of congestion detection and response algorithms [Stevens] that have evolved over the past 10-15 years in the Internet environment, and in that environment they work very well. (The Internet operates under conditions of very high load and high potential congestion, but congestion is generally avoided and the links usually operate at high utilization efficiency.) The problem in ATM networks is that the TCP control algorithms are based on certain assumptions about the relationships among parameters like the packet switch buffer sizes, link minimum transmission unit sizes, link bandwidth, etc. It is very easy for ATM networks to violate these assumptions, and when this happens both the throughput and the link utilization efficiency can be very low in the presence of congestion. (In current implementations of ATM networks, this congestion is frequently caused by small output port buffers in ATM switches. Fortunately, most new switches are being designed with large output port buffers.) For a detailed analysis and case study of this situation, the reader is referred to "Performance Analysis in High-Speed Wide-Area ATM Networks: Top-to-Bottom End-to-End Monitoring" ([Tierney]).

3.5 Security Architectures for Large-Scale Remote Environments

The access to remote resources via open networks, whether for monitoring, control of experiments, or manipulation of data, requires that authentication and access control mechanisms be in place in order to provide safeguards against unauthorized access to resources, and to assure privacy of proprietary, confidential, or otherwise restricted data. Therefore, awkward, intrusive, and expensive security will delay wide-spread use of remote resources in open environments.

In addition, we would like the security architecture and supporting infrastructure to be general enough to enable automatic brokering of resources. This new capability could support a scenario where resource owners (either as a mainline business, or as a barter-based use of excess capacity) advertise resources and user agents collect and assemble these resources into useful systems. (This could be a very useful service in the world of scientific experiment control, when significant resources may be required only for well-defined periods of time, and idle otherwise.) It is our hypothesis that this capability can be built on the same use-condition based security architecture that is described here.

In order to meet the need (that is, in order to provide the required capabilities in a manageable, cost effective, and usable manner) the security architecture must be flexible and effective, and easily deployed, administered, and used. In order to realize its potential not only for transparent access control, but also as the basis of an automated resource brokering system, all aspects of the security process after certificates are obtained must be automated.

The problem, then, is to identify and demonstrate authentication and access control methodologies that are applicable to all user-level resources - data, storage and processing, control of external devices, etc. - in widely distributed systems, and that are also general, scalable, "strong", non-intrusive (for both the administrator and the user), and easily managed.

The overall goal of the security architecture is to be able to encode, distribute, protect, and then act on, information that is needed to provide routine, secure availability of remote resources in a widely distributed environment, and to enable a level of automated processing that will permit resource brokering.

For more information, see [Johnston96a].

3.5.1 Background

Public-key certificates (PKC) are an inherently reliable and secure mechanism for widely distributing the assured information needed for authentication and authorization. (See [RSA] for background information.)

Public-key certificates can encode information about a principal, or information expressed by a principal, or a relationship between principals, in a secure and verifiable way. Certificates that provide some policy based assurances of the identity of the principal we call identity certificates. Certificates that encode organizational / group affiliation, creditworthiness, level and scope of training, etc., we call “credentials”. Certificates that encode use-conditions, e.g., cost, required role attributes (personal identity, group membership, organizational function), required personal characteristics (training, credit worthiness, etc.) we call requirement certificates. Certificates that encode a relationship of a principle to a policy (e.g. that one certification authority operates under the policy of another) we call trust certificates.

There are several infrastructure elements that are needed to support our model.

There must be a mechanism to establish and represent trust relationships (policy and procedure agreements) among authorizing or verifying entities. The common ways of doing this are represented in a continuum from a centralized root, hierarchical structure of certification authorities (CA) at one end of the spectrum (see [Kent]), through “webs” of small (organization-scope) CAs, to the completely decentralized approach of PGP (where individuals attest to each other’s attributes) at the other end. All of these approaches are in use, and we have focused on the middle ground of organization-level CAs that attest to specific facts and establish trust on a pair-wise basis among themselves.

There must be mechanisms for making certificates locatable based on content, and widely and reliably available. There are several current approaches to this. X.500 directory servers provide, in principle, a standard way of searching for and distributing X.509 identity certificates. There are already WWW sites and ftp sites that provide PGP certificates. Other distributed information mechanisms (e.g., whois++) are possible. Another approach is to use Web-based searching of textual representations of certificates. There are very fast search engines for text searching, and techniques to search the distributed environment over a large space of textual data sets. The results of a search would be URL pointers to the certificates that match the search criteria. If the certificates are binary objects, then this approach is essentially identical to an image library capability for indexing large data-objects. (See [ImgLib] and [Johnston95b].)

Our current approach is focused on using the “Lightweight Directory Access Protocol” [LDAP] to communicate with X.500 servers to obtain X.509 identity certificates, and on the Web search method mentioned above for credentials and requirement certificates.

There must be a machine-comprehensible encoding for the representation and automatic manipulation of use-conditions and attributes. While our first implementation will use ad-hoc representations, we are also looking at approaches like the DCE “Authorization for Distributed Applications and Groups” (ADAGE) project (see [ADAGE]) and the generalized certificates being discussed in the IETF SKPI working group [SPKI].

3.5.2 A Use-Condition Centered Security Model

We are building a use-condition centered model that is based on verifiable and securely represented resource controller / owner imposed use-conditions, and equally secure and verifiable satisfaction of those conditions. The individual identity that is so important in many security architectures is just one of many factors in our model.

The basic idea is that the principals that control resources will establish a set of conditions for the use of the resources. The use-conditions are whatever is appropriate to the resource: a level of training for operating an instrument; membership in a group (e.g. health care professionals working for a particular HMO) for access to patient data; creditworthiness for access to a resource that requires a commitment of payment; etc. Use-conditions are encoded in signed certificates and deposited in a “public” server.

Paired with the resource owner’s use-conditions are the principals that can attest to the relevant attributes of a user or agent that is seeking access to a resource: A school registrar may certify a level of training; a state examining board may certify a level of professional attainment; an organization can certify employment (and related group membership); a bank, a level of credit; and a certification authority may attest to individual identity. Again, the principal (person or organization) that can attest to some attribute, expresses that fact in a signed certificate that is deposited in a “public” server. See Figure 11.

In addition to use-conditions and corresponding credentials (certification of an attribute that satisfies a use-condition), a collection of trust relationships must exist. So, for example, for a resource controller to accept a certificate that purports to represent satisfaction of a use-condition, the resource owner must trust that the principal that is attesting to (certifying) an attribute of the user, is qualified and/or certified to make such a claim. These trust relationships may be direct: the resource owner and the verifier of conditions may have an explicit, pair-wise agreement; or they may be indirect: a resource owner may accept a statement of policy, perhaps policed by a third party, as an implicit guarantee of the satisfaction of a use-condition.

To gain access to network-based resources, a broker / agent, acting on behalf of the potential user, collects the certified use-conditions for a resource, and then collects the credentials that verify the user’s satisfaction of the use-conditions. The combination of resource identity, use-conditions, and matching credentials, each of which is cryptographically secured by the relevant principal, are used to form a “capability”. A capability allows the user to access, monitor, manipulate, use, etc., a remote resource.

In order to enable building systems dynamically from network resources, a similar procedure could be used. As with access control a broker / agent, acting on behalf of the potential user, collects the certified use-conditions for a resource, and then collects the credentials that verify the user’s satisfaction of the use-conditions. In addition to use-conditions we will assume a payment requirement. One way of dealing with this is to use a “cybercash” approach, where the cash or coupons are denominated in appropriate units: monetary, resource units, barter units (resource units for other resources), etc. These coupons are presented along with the capability, but are not “consumed” until the resource access is actually accomplished. One could easily imagine coupling this into a commercial banking / settlements system to handle accounting.

The last part of the model are “deferred” verifications. These actions may be required in the use-conditions, but are deferred to the resource access control mechanism. In a sense, the consuming of a coupon is an example of this. Other examples include checking all manner of time sensitive

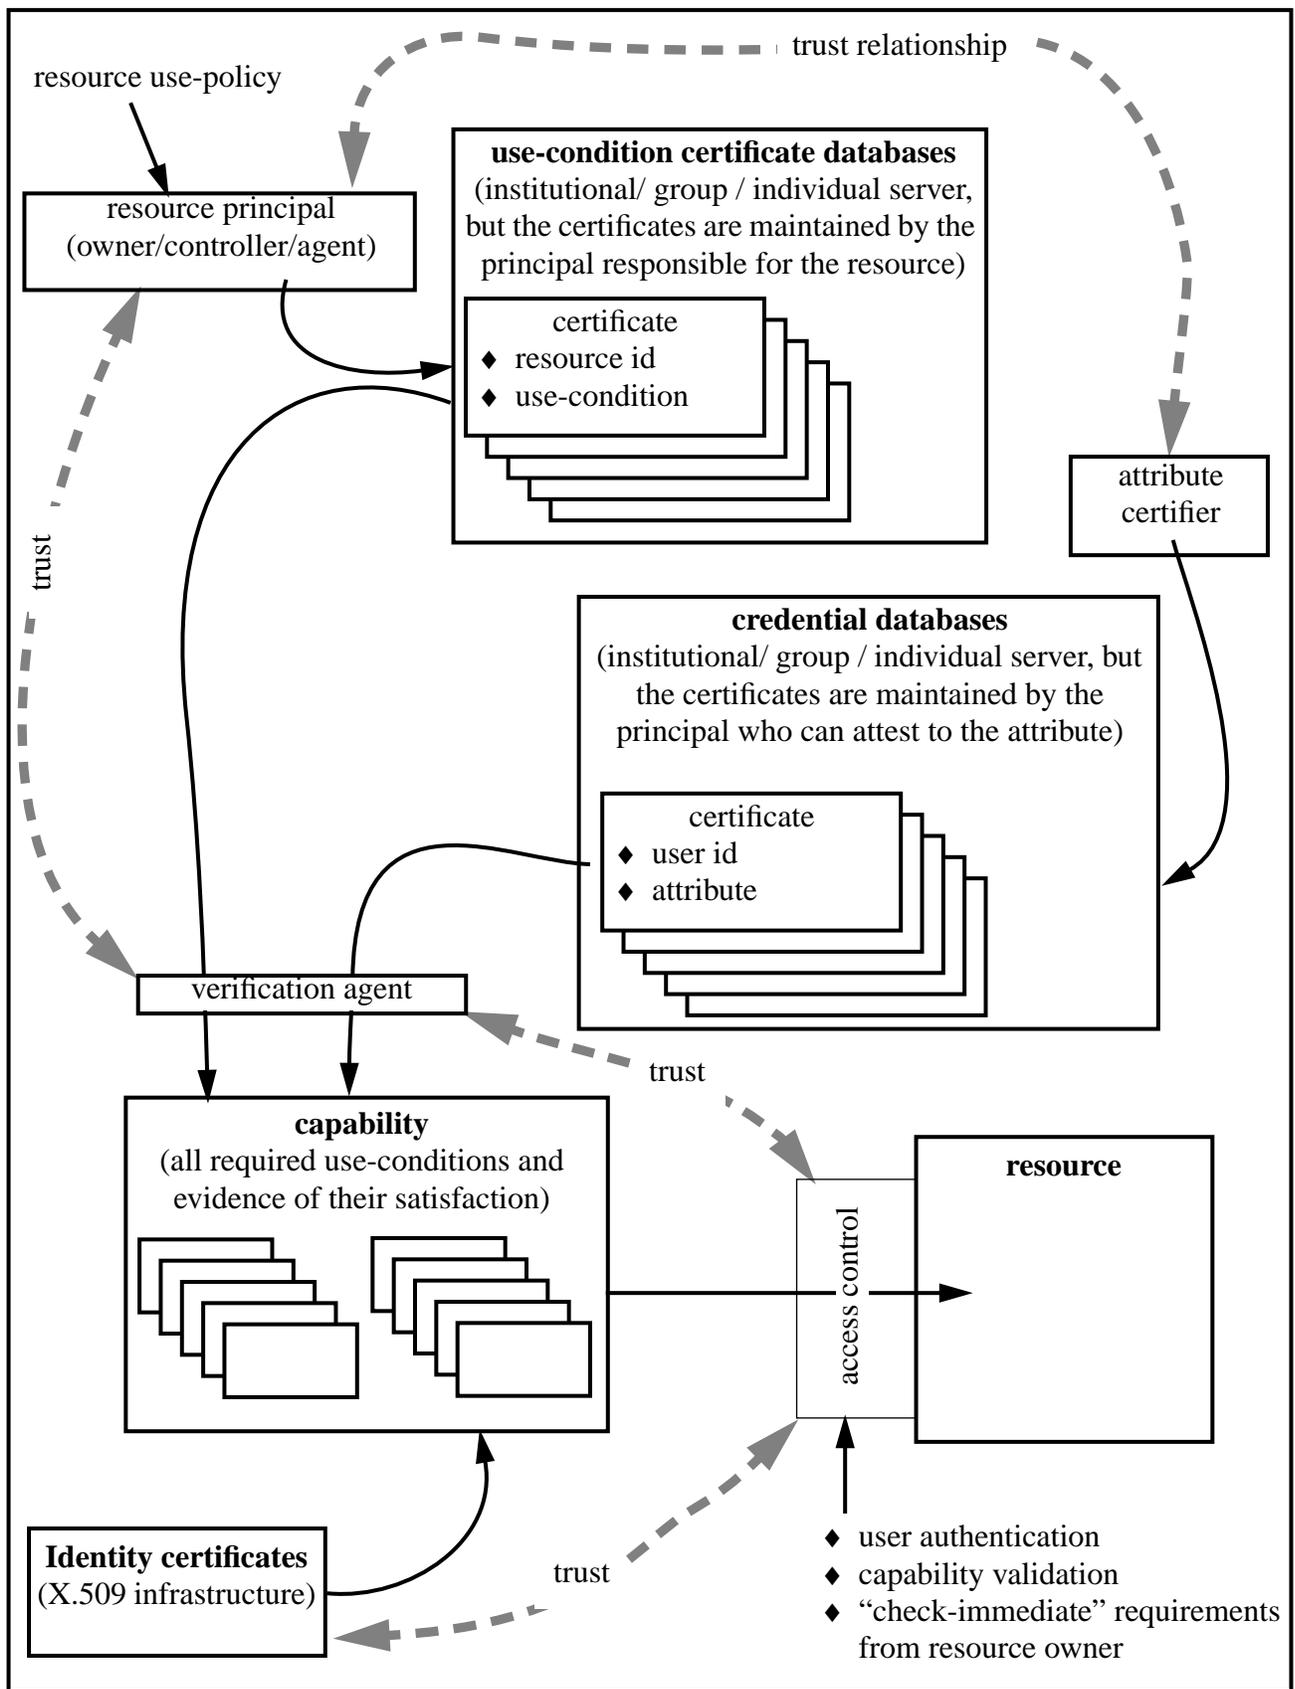


Figure 11

The Security Model

requirements, re-verification of user credentials in the case of high-consequence access, etc. In many - perhaps most - cases, presentation of a valid capability - i.e. all use-conditions matched by corresponding credentials, etc. - would be all that is required for access to the resource. (That is, the resource controller would only have to authenticate the user and the signer of the capability (assuming that a trusted third party originally assembled and verified the capability). "Check immediate" requirements beyond this would be the result of specific policy requirements of the resource principal.

3.5.2.1 Basis in Practice

Our general approach is motivated by the credential model used in most Western societies for consumer contracts. An example that illustrates many of the points is that of renting an automobile at an airport. The essential elements of this process include presenting and verifying a set of credentials, and then being given a capability (authorization certificate) in the form of a validated rental car contract, that is then used to gain access to the resource (the automobile).

In this model, a driving license is an attribute certificate that represents a guarantee by a trusted third party (the issuing state Dept. of Motor Vehicles) that an individual has had appropriate training and is currently authorized to drive a car. The driving license also plays the role of identity certificate when the individual's signature must be verified in order to signify agreement to a contract. An individual's physical signature is a "private" or individual expression (like the private-key part of the public-private key pair). The photograph of the individual's face, and the photo of a written signature together on a driving license certificate constitutes the certified "public" expression of identity.

One's credit card represents a creditworthiness credential, and provides the entree for on-line revalidation of creditworthiness. These two credentials (driving license and credit card) are verified (usually automatically), and a capability - to use the rental car, with restrictions - are represented in a third, combined authorization certificate (or "capability"): the rental car agreement.

The rental car company agents are the verification agents. They verify multiple use-conditions by checking a collection of attribute certificates and then issue a (composite) authorization certificate. We call this new certificate a "capability" because it associates a set of otherwise unrelated certificates to represent the satisfaction of a set of use-conditions that, together, enable the use of a specific resource. The guard at the rental car lot gate is the resource access controller. He checks that the capability is signed by the verification agent and authenticates the user (checks driving license again). The capability, then enables the use of the rental car (the resource).

The success of this model is clear: most responsible people can rent an automobile - a dangerous and expensive resource - anywhere in the world by producing the commonly possessed set of credentials: the driving license and the credit card to satisfy the use-conditions imposed by the resource owner.

This example embodies most, if not all, of the characteristics of a security architecture that can provide secured global capabilities in an open Internet and the abstraction of these operations are necessary, and probably sufficient, conditions to ensure general resource security in the Internet.

3.5.2.2 Strengths of the Model

One of the expected advantages of the model described here is that it should provide a natural (and distributed) responsibility for information validity, placing the basic expression of policy (issuance of a use-condition certificate or an attribute credential) with those principals that have the direct responsibility for the information (e.g., the home state Department of Motor Vehicles) or for setting use-conditions (e.g., the rental car company).

The combination of direct reference to the principals responsible for the generation and assurance of information (the resource owners or policy makers), and standardization of representations of the information that permit automated verification, should form the basis of an easily used, scalable, and general purpose authentication and authorization system that provides such services as qualifying remote experimenters to use instrumentation systems in distributed collaboratories, as well as forming the basis for dynamically configured, large-scale, distributed computing resources.

Automated acquisition and verification of attribute credentials will increase efficiency and reliability, and will play a critical role in scaling remote environments by removing the need for human intervention in repetitive and trivial tasks such as credential checking, especially when the human participants are never in the same place at the same time, may never physically visit the site where service is being requested, and may not have any knowledge of each other outside of the contents of the credentials.

However, even with the tools and infrastructure in place, they are still just that: tools and infrastructure. In each given situation where a resource is being protected, those responsible must evolve a security policy model that defines the level of assurances, protection, auditing, etc., that meets their requirements. Flexible tools and infrastructure make it possible to implement specific security policy in a usable and cost effective way.

3.5.2.3 Summary

Our approach to meeting the goals of security for remote access to resources on open networks has several aspects.

Decentralization: To provide security in large-scale, open access environments, and to meet the objectives for deployability, we have evolved a decentralized model that is based on use-conditions imposed by the principal responsible for a resource. That is, while the access restrictions are enforced at the resource, the conditions of access are specified by the principal responsible for the access or use policy for the resource. The responsible principal (or its agent) is the only one who imposes and records use-conditions, and those conditions are made available in a way that has nothing to do with the operation of the resource itself. The security mechanisms at the resource are only responsible for providing access control that ensures that the use-conditions are satisfied. To do this a collection of verified use-conditions and credentials are presented as a capability. The verification is performed by a third party that is trusted by the resource owner to determine that all of the use-conditions are satisfied by the presented credentials. Finally, the access control agent performs a set of “last minute” (deferred) checks: authenticating the user presenting the capability, authenticating the verification agent that signed the capability, and performing any deferred actions required by the resource owner (e.g. credential revalidation, payment, etc.).

As with many scalable distributed systems, the decentralized aspect of the model is a key to scalability: the administrative task of maintaining an appropriate expression of access policy is only done by those entities (principals, or their agents) that are directly responsible for the policy governing the use of a resource. The operational administration of the resource has only to deal with providing access control mechanisms, not with access control information, which is presented (both use-conditions and their satisfaction) in order to gain access the resource. The access control mechanism also has to deal with “deferred” requirements. Central administration of information like access control lists, or multiple roles that have to be remembered by resource users, are not scalable approaches in a global distributed environment, and are aspects of security that this model is intended to replace.

We believe that the combination of this decentralized model, together with the generalized public-key certificate infrastructure that supports it, can provide a highly scalable and transparent infrastructure for both authentication of, and authorization for, access to the full range of network based remote resources.

Automated processing: Two main features will be needed to enable automatic brokering. One is an approach to encoding the relevant information (both use-conditions and attributes that can satisfy them) in ways that can be parsed, understood, and then manipulated by software systems. The second is widely available (e.g. network-based), searchable databases that can be used to “advertise” resources and use-conditions, and to provide repositories for the certificates that can attest to the satisfaction of use-conditions.

Deployable: To address deployability, we have evolved the decentralized architecture noted above, and an implementation design that manages certificates on behalf of the users, should provide “strong” security for applications that use the security architecture directly, and should provide “good” security for as many unmodified applications as possible.

3.5.3 Implementation of a Prototype Architecture

Our initial prototype addresses:

- User authentication, credential, and key management via SSH [SSH]
- Identity certificate management via LDAP/X.500 [LDAP] / WWW
- Resource access control
 - security context establishment via GSS/SPKM [Adams]
 - secure, authenticated, and integrity protected exchanges via GSS-API [GSS]

The SSH design provides a convenient and transparent management of the credentials and public keys of servers and the private keys of the user. (Roles are possible through different “identities” of the user.) We have modified the SSH credential and key management facility to be able to set up the SPKM security context for GSS. (See [Adams] for a good discussion of the rationale for SPKM and its relationship to GSS/Kerberos.)

SSH makes the user’s private key available to SPKM for user authentication and to do the session key exchange needed for GSS security context establishment. Once this is done, the GSS-API provides an application interface to send validated and/or confidential communication.

GSS-API (Generic Security Service - IETF Common Authentication Technology working group) provides a simple application interface to send validated and/or confidential communication:

- security context establishment
 - get your credentials from ssh-agent
 - get server (resource provider) public key
 - do mutual authentication
- very simple application interface
 - GSS_sign (context, message)
 - GSS_seal (context, message)
 - (GSS_verify, GSS_unseal)

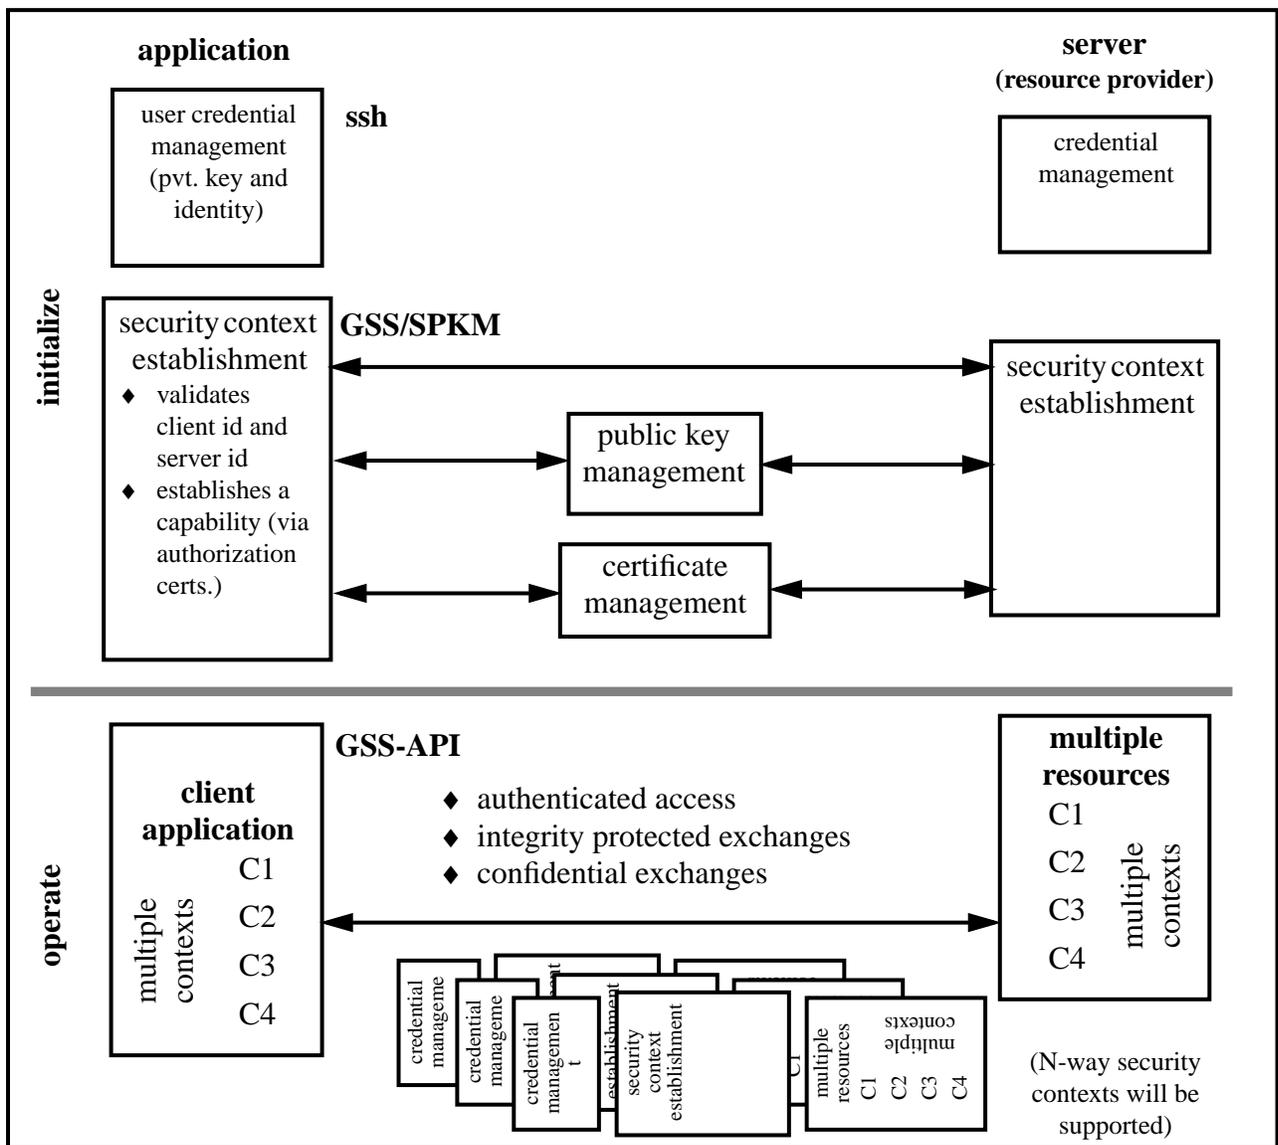


Figure 12 Current Prototype Security Architecture Implementation

The current prototype can be represented as in Figure 12.

For search and recovery of identity certificates we are examining the LDAP and Web search engine approaches mentioned above.

We currently use an ad-hoc use-condition encoding, but are examining other approaches like ADAGE, noted above. The automatic brokering depends on being able to collect, parse, and automatically match use-conditions and credentials. This is future work.

3.5.4 Relationship to Other Work

Our general approach is an application of public key certificates, and as such it draws ideas from an active community of researchers.

The basic concepts of distributed systems security are laid out in various ISO (International Organization for Standardization) and ECMA (a Europe based international association for the

standardization of information and communication systems) documents (see [SIRENE]). SESAME is an implementation that incorporates many of the [ECMA-219] and [ECMA-138] concepts. (See [SESAME].)

We have, however, taken a somewhat different direction than the ECMA and SESAME approaches with respect to attribute certificates, which are central to our approach. In the ECMA approach, most of the emphasis on the use of attribute certificates seems to be in the form of access control lists and roles (party A, who controls a resource, grants to party B the right to use that resource, perhaps under restricted conditions (roles)). Our approach treats use-conditions, rather than identity, as the fundamental concept. That is, we see resources primarily being protected by requirements on the user (e.g. level of training) rather than identity. (ACLs are, of course, easily implemented in this model.)

Our implementation is based on some refinements to the “secure shell” [SSH], combined with the “Generic Security Service Application Program Interface” [GSS] to provide an easily used application security mechanism.

Although evolved independently, the idea of CAs dealing only with identity certificates and resource principals signing use-condition certificates, is very similar to the generalized certificate ideas described by Carl Ellison [CME].

3.5.5 Prototype Application

We are using the prototype implementation to provide security to an experimental health care information application (see [Kaiser-CalREN]) that uses the Distributed-Parallel Storage System ([DPSS]). The application uses a WWW-based system for managing certain kinds of medical imaging records. The data-object viewing application (“Vplayer”) is invoked via a WWW image database frontend that uses WWW security to protect the metadata. The application (a specialized, combined video and image browser) then accesses data stored on the DPSS. The Vplayer access to DPSS data is protected by the security architecture described here. (That is, credentialing is handled by the secure shell that runs Vplayer, and the DPSS access is controlled by checking authorization certificates presented by Vplayer.)

4.0 Conclusions

We have described several elements of an architecture the using high-speed, wide-area networks as the basis of a highly distributed environment for handling all aspects of large data-objects. Variations of the overall approach have been applied in several different projects, and we have described some of the results, in particular as they apply to a distributed health care information system.

We believe this approach to be both practical and powerful given the increasing availability of high-speed networks. Further work will be done on both the various functional elements and specific applications. Please see <http://www-itg.lbl.gov> for more information.

5.0 Acknowledgments

The overall distributed large data-object approach has evolved over several years with funding from both the Mathematical, Information, and Computational Sciences Division of the Office of Energy Research, U. S. Department of Energy, and ARPA Information Technology Office, both of which are gratefully acknowledged.

In the Kaiser CalREN project, E-J. Pol of Philips Research, Palo Alto, CA, designed and implemented the workstation interface for the cardio-angiography system. Dr.J. Terdiman, of the Kaiser Permanente Division of Research provided the leadership for the health care aspects of the project. Dr. Bob Lundstrum, Kaiser San Francisco Hospital, Cardiac Catheterization Laboratory, also provided guidance and information specific to the health care application. Pacific Bell, through its CalREN program, provided funding for connecting the LBNL, Kaiser, and Philips sites to the SF Bay Area ATM network. Sun Microsystems assisted in various ways, including loaning equipment to several Kaiser sites.

Finally, one of us (Johnston) would like to acknowledge Ace Allen, M.D., of the KU Medical Center in Kansas City for enthusiastic early support of the idea of networked medical information systems; Ms. Sue Kwentus, RN, and the MCC HOST/OSL group for focusing some of early ideas; and Sprint generally, and particularly Lt. Col. John Strand, U.S. Army (ret.), now Director of Technology Planning and Integration at Sprint, for support of the MAGIC backbone network (without which this work probably would not have been done).

6.0 Notes and References

ADAGE “Authorization for Distributed Applications and Groups”, see <http://www.osf.org/www/adage/index.html>

Adams Adams, C. “IDUP and SPKM: Developing Public-Key-Based APIs and Mechanisms for Communication Security Services”. See <http://bilbo.isu.edu/sndss/sndss96.html>

Blaze “Blaze, M., J. Feigenbaum, J. Lacy, “Decentralized Trust Management”, to be presented at the 1996 IEEE Symposium on Research in Security and Privacy, Oakland, CA, May 1996. (<ftp://research.att.com/dist/mab/policymaker.ps>)

BAGNet-1 “Bay Area Gigabit Testbed (BAGNet)” See <http://www-itg.lbl.gov/BAGNet.html>

BAGNet-2 Pacific Bell Announces First CalREN Award” <http://www-itg.lbl.gov/BAGNet.hm.pg.docs/PacBell.BAGigabit.Press.Release.html>

CME Ellison, C., “Generalized Certificates”, <http://www.clark.net/pub/cme/html/cert.html>

CRYPTOLOG ““The Internet Guide to Cryptography”” <http://www.enter.net/~chronos/cryptolog1.html>

DPSS The Distributed-Parallel Storage System (DPSS) Home Page (<http://www-itg.lbl.gov/DPSS>)

ECMA-138 “Security in Open Systems - Data Elements and Service Definitions”, 1st Edition (December, 1989). European Computer Manufacturers Association, <http://www.ecma.ch/ecma-138.HTM>

ECMA-219 “Authentication and Privilege Attribute Security Application with related key distribution functions”, 1st Edition (December 1994). European Computer Manufacturers Association, <http://www.ecma.ch/ecma-219.HTM>

GSS Linn, J., “Generic Security Service Application Program Interface” (<http://ds.internic.net/rfc/rfc1508.txt>) Also see more recent and related drafts at the IETF Common Authentication Technology home page (<http://www.ietf.cnri.reston.va.us/html.charters/cat-charter.html>) and at <http://www.ietf.cnri.reston.va.us/ids.by.wg/cat.html>.

Hudgins-Bonafield "Pipe Dreams And Lessons With TCP Over ATM". Published in Network Computing Online (<http://techweb.cmp.com:80/techweb/nc/docs/default.html>). See <http://techweb.cmp.com:80/techweb/nc/706/706hreportb.html>

ImgLib See http://www-itg.lbl.gov/ImgLib/ImgLib_intro.html

Johnston96a Johnston, W. and C. Larsen, "Security Architectures for Large-Scale Remote Collaboratory Environments: A Use-Condition Centered Approach to Authenticated Global Capabilities" (draft at <http://www-itg.lbl.gov/~johnston>)

Johnston95a Johnston, W., J. Terdiman, E. J. Pol, "The Use of Public ATM Networks for Health Care Imaging Information" (see <http://www-itg.lbl.gov/Kaiser/LKP>)

Johnston95b W. Johnston, and D. Agarwal, "The Virtual Laboratory: Using Networks to Enable Widely Distributed Collaboratory Science" A NSF Workshop Virtual Laboratory whitepaper. (See <http://www-itg.lbl.gov/~johnston/Virtual.Labs.html>)

Johnston95c Johnston, W., "BAGNet: A High-Speed, Metropolitan Area, IP over ATM Network Testbed", IEEE Compton'95, March, 1995, San Francisco, CA. Also see <http://www-itg.lbl.gov/BAGNet.hm.pg.docs/CompCon.95.3.html> .

Johnston93 W. Johnston and A. Allen, "Regional Health Care Information Systems: Motivation, Architecture, and Implementation," LBL Technical Report 34770 (draft), Dec. 1993.

Kent Steve Kent, "Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", RFC-1422, <http://ds.internic.net/rfc/rfc1422.txt>.

Kaiser-CalREN Kaiser - LBNL - Philips CalREN project. See <http://www-itg.lbl.gov/Kaiser/LKP/homepage.html> .

LDAP "Lightweight Directory Access Protocol", <http://www.umich.edu/~rsug/ldap/ldap.html>.

MAGIC testbed, <http://>

Mendez and Huitema "A New Approach to the X.509 Framework: Allowing a Global Authentication Infrastructure without a Global Trust Model", S. Mendez and C. Huitema, in "Proceedings: Symposium on Network and Distributed System Security", IEEE, San Diego, CA, February 16-17, 1995.

MC "Visa & MasterCard Combine Secure Specifications For Card Transactions On The Internet Into One Standard" <http://www.mastercard.com/Press/release-960201.htm>

NTON The National Transparent Optical Network (<http://www.cais.net/nton/>) connects a number of organizations in the San Francisco Bay Area)

RSA "RSA Labs' Frequently Asked Questions About Today's Cryptography" http://www.rsa.com/rsalabs/faq/faq_home.html

SESAME

"SESAME (a Secure European System for Applications in a Multi-vendor Environment) is a European research and development project, part funded by the European Commission under its RACE programme. It is also the name of the technology that came out of that project.

The SESAME technology offers sophisticated single sign-on with added distributed access control features and cryptographic protection of interchanged data.

SESAME is a construction kit. It is a set of security infrastructure components for product developers. It provides the underlying bedrock upon which full managed single sign-on products can be built.”

See <http://www.esat.kuleuven.ac.be/cosic/sesame3.html> .

SET “Secure Electronic Transactions” <http://www.mastercard.com/set/set.htm>

SIRENE “Security in Computer Networks” (<http://www.zurich.ibm.com/pub/sti/www/g-kk/sirene/>
also see <http://www.zurich.ibm.com/pub/sti/www/g-kk/sirene/pointers.html>)

SPKI “Another public-key infrastructure working group has started in the IETF (SPKI). This group has formed to explore public-key certificates that are attribute-based instead of name-based, for use in a variety of Internet applications.” Subscriptions: majordomo@c2.org; send “subscribe spki”

Spartan Sprint test network. T-3 / OC-3 ATM from Kansas City, Kansas to the Sprint Test Laboratory in Burlingame, CA.

SSH Ylonen, T., “The SSH (Secure Shell) Remote Login Protocol” (<ftp://ietf.cnri.reston.va.us/internet-drafts/draft-ylonen-ssh-protocol-00.txt>). Also see “Ssh (Secure Shell) Remote Login Program” (<http://www.cs.hut.fi/ssh/>).

SSLeay Young, E. “SSLeay”. (SSLeay is a free implementation of Netscape’s Secure Socket Layer - the software encryption protocol behind the Netsite Secure Server and the Netscape Browser.)

See <http://www.psy.uq.edu.au:8080/~ftp/Crypto> and <http://www.homeport.org/~adam/crypto/ssl.html>

“What started as an effort to implement the SSL protocol has turned into a fairly complete cryptographic library (which Eric is still working on). There is also quite a bit of ASN.1 support, with routines to convert and manipulate the base ASN.1 types, X509v3 certificates, certificate requests, certificate revocation lists (CRL), RSA private keys and DH parameters. There are routines to load and write these objects in base64 encoding and routines to convert ASN.1 object identifiers to/from ASCII representations and an internal form. There are functions for verification of X509 certificates and for specifying where to look for certificates to ‘climb’ the x509 ‘tree’. This last part of the library is still evolving.

The big number library is quite complete and has no restrictions on the size of the numbers manipulated. RSA and Diffie-Hellman routines have been layered on top of this library.”

Stevens Stevens, W. “TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms”. IETF Draft. See <ftp://ietf.cnri.reston.va.us/internet-drafts/draft-stevens-tcpca-spec-01.txt> .

TerraVision A terrain visualization application that uses the DPSS to supply high-speed data streams for realtime, “virtual reality” type navigation of real landscape. (See <http://www.ai.sri.com/~magic/terravision.html>)

Tierney Tierney, B., W. Johnston, G. Hoo, J. Lee, “Performance Analysis in High-Speed Wide-Area ATM Networks: Top-to-Bottom End-to-End Monitoring”, To be published in IEEE Networking Journal, LBL Report 38246, 1996. (Also see <http://www-itg.lbl.gov/DPSS/papers.html> .)

Wiltzius Wiltzius, D., L. Berc, and S. Devadhar, “BAGNet: Experiences with an ATM metropolitan-area network”, *ConneXions*, Volume 10, No. 3, March 1996. Also see <http://www.llnl.gov/bagnet/article.html> .