B. Executive Summary

B.1 Synopsis

Audio and video data forms a significant and growing fraction of the total information being created and archived. The diverse nature of this data combined with its magnitude poses difficult problems in organizing, storing, and delivering it in a secure and timely manner. As multimedia information plays an increasingly important role, efforts to solve these problems are becoming critical. Networked multimedia information systems are viewed as catalysts for new research in many areas of computer science and engineering, ranging from basic research to applied technology and education. This view is a result of the fact that no single monolithic service architecture can meet the wide spectrum of characteristics and requirements of various multimedia applications. One size does not fit all in this medium of communication.

Critical issues of managing and translating integrated quality of service requirements across distributed multimedia systems form the unifying theme of the proposed research infrastructure. Strongly coupled to this theme are issues relating to specification, support, and provisions for required levels of quality of service via distributed multimedia servers. When viewed in conjunction with real world constraints and system-wide end-to-end performance requirements, these present formidable research and implementation challenges that encompass all the components of a multimedia support infrastructure (MSI).

The ultimate objective of a comprehensive end-to-end quality-of-service specification relies on several individual component systems such as networks, databases, and security. This research project aims to investigate each of these components with a view to optimizing overall performance. The heterogeneous notions of quality embodied by various components in the MSI infrastructure raise the following research issues:

1. How do we support end-to-end quality of service across capture, storage, delivery, and presentation of multimedia data?
2. How do we effectively translate and inter-operate quality-of-service requirements across subsystem boundaries? This support includes, for example, mapping user level quality of presentation (QoP) to supported quality of service (QoS) at all underlying levels of the system.

The proposed research and development related to these issues will be undertaken with a set of sample applications from Nuclear Engineering, Veterinary Medicine, and a Purdue facility for Web-based distance learning called Purdue-Online. This project will focus on the following three principal areas: networks, databases, and security. It aims to investigate the contributions of each of these components to overall system performance as they relate to the aforementioned research issues. The research will combine theory with extensive experimentation.

Figure 1 shows the overall structure of the research project and how it relates to the motivating applications and requested infrastructure.

The proposed infrastructure will support the following tasks:

1. *Multimedia research testbed construction*. In order to accomplish the tasks outlined in this proposal, access to state-of-the-art hardware is essential. We propose to acquire the equipment necessary to create a world class multimedia computing environment to be used as a multimedia support infrastructure (MSI). The architecture and configuration of this testbed MSI will form the substrate for our research team. (Task leader: Korb. Participants: Elmagarmid, Park, Tsoukalas, Ishii, and Coppoc).
2. *QoS management for networked multimedia*. We propose to use the MSI infrastructure to perform research in QoS-sensitive access, dissemination, and transport of multimedia data retrieved from distributed multimedia databases. (Task leader: Park. Participants: Yau and Bhargava)
3. *Distributed multimedia database management*. Systems of the future will be based on highly connected information domains and multimedia databases. We propose to employ the MSI infrastructure to perform research in distributed multimedia databases. Particular target research issues include: quality of presentation management, specification, and translation as they relate to: data storage, I/O management, data placement, meta-schema design, data replication for reliability and efficiency, admission control for user sessions, and benchmarking of distributed applications. We will also develop and benchmark

multimedia compression techniques for the transport of complex multimedia objects across the distributed networks. (Task leaders: Elmagarmid and Ghafoor. Participants: Delp, Bouman, and Prabhakar)

4. *Security at user and network levels*. Security is a major concern for information systems that may be deployed in practice. We are planning to use the MSI infrastructure to perform research in ATM security, audit trails, watermarking, and intrusion detection of multimedia networks and data. (Task Leader: Spafford. Participants: Atallah, Wagstaff, Delp, and Park)

5. *Multimedia capture and presentation*. We plan to use the MSI infrastructure for realistic capture, compression, and presentation of multimedia data. To this end, we plan to employ real-world multimedia data applications already accessible to us. The research in this category will be limited to compression. The rest of the issues in capture and presentation are carried out using commercial off-the-shelf technology. (Task leader: Delp. Participants: Szpankowski, Grama, and Nour)

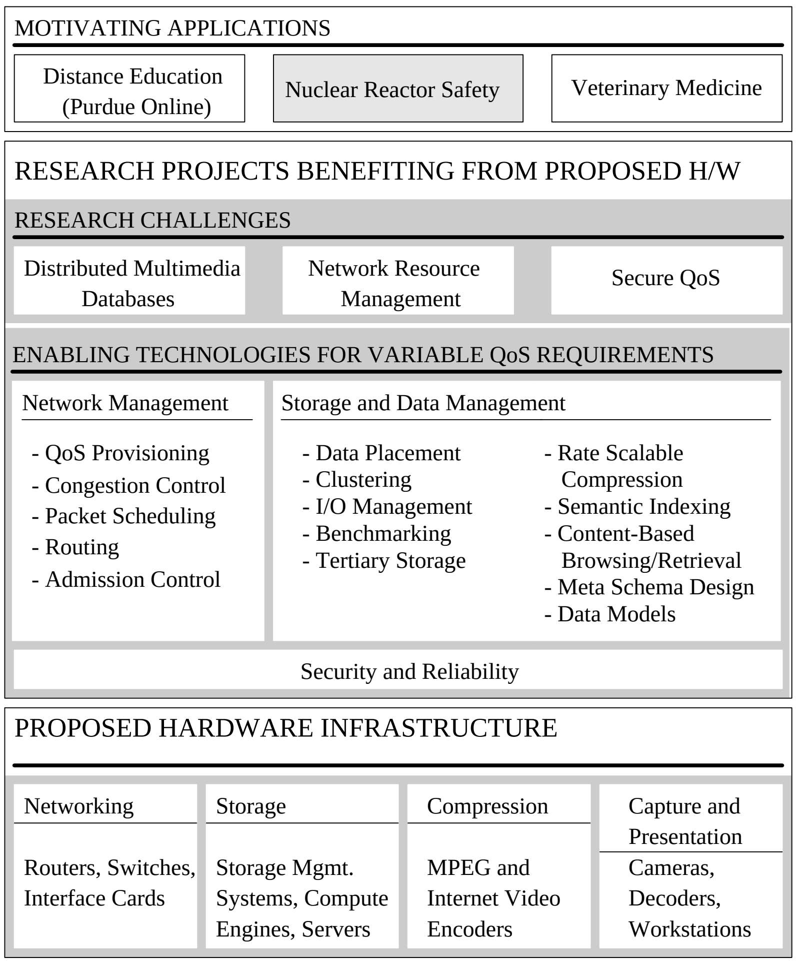

Figure 1: Applications, Challenges, and Infrastructure of the Project

B.2. Rationale

The team collaborating on this NSF infrastructure project includes 18 faculty members from four departments. The proposed project is expected to enhance the high level of synergy that already exists among the participating research teams. It is the intention of this project to bring together researchers in the fields of databases, networks, computer security, image processing, nuclear engineering, and veterinary medicine. The team also involves a mix of faculty in various stages of their career development. While many of these researchers already have joint research projects and funding, the proposed infrastructure will help create a common testbed to which they will all contribute. To this end, we will initiate a monthly seminar on multimedia computing to foster a greater degree of interaction between researchers.

The proposed infrastructure along with the research projects supported by it will have a significant impact on the quality of instruction in networks, security, and databases. The infrastructure will also be available to researchers from Fisk University, a historically black university and our partner in a recent award from EPRI/DoD. Researchers from Fisk will be able to fully use the infrastructure for their own work as well as our joint funded research. They will be encouraged to spend summers at Purdue along with their graduate students to get involved in this project.

Purdue has been successful with its women-in-engineering programs. Last year, 25% of the entering freshmen engineering students were women. The retention rate is very good: Women comprise 25% of the graduating class. And women tend to do better academically, as a group, than men. We are proud of Purdue's record but we can do better. We want to engage through REUs a larger number of female students in the use of this proposed world class MSI facility.

The central equipment that will be available for use by all the researchers will be housed in the MSI Laboratory (room Math 431). The rest of the equipment will be distributed among all the labs in which the research will take place. The mix of networking, database, and compression hardware we are requesting will create a world class infrastructure that complements what we already have.

We have been able to obtain substantial commitments from Industry and the University in terms of specific new funds to support this project far beyond the 33% requested by NSF. The following table summarizes the proposed budget. Out of the estimated total cost of this project of $2,201,212, we request $1,386,613 from the NSF. This constitutes a cost sharing of 37%. The industry cost share from Bellcore is in the form of support for one research assistant for the five years of the project. A license agreement for the CineBlitz media server from Lucent Inc. is awarded at no cost to Purdue. The combined existing research funding of the project team is over $10 million.

Five Year Summary Budget (in units of $1000)

	Purdue	**Industry**	**NSF**	**Total**
Personnel	$674	$141	$0	$815
General Expenses	$0	$0	$420	$420
Equipment	$0	$0	$966	$966
Total	$ 674	141	$1,386	$2,201

C. Research Infrastructure Description

The research infrastructure we are building, shown schematically in Figure 2, is designed to accommodate and advance research in the area of quality of service networking for multimedia databases. Specifically, we are creating a testbed to build and experiment with algorithms and techniques to compress, store, index, retrieve, and deliver multimedia data to applications. Purdue faculty associated with this proposal are actively involved in research in three of the areas shown in this figure: quality of service networking, multimedia databases (storage, indexing, and retrieval), and compression. In addition, the security research pervades the entire infrastructure.

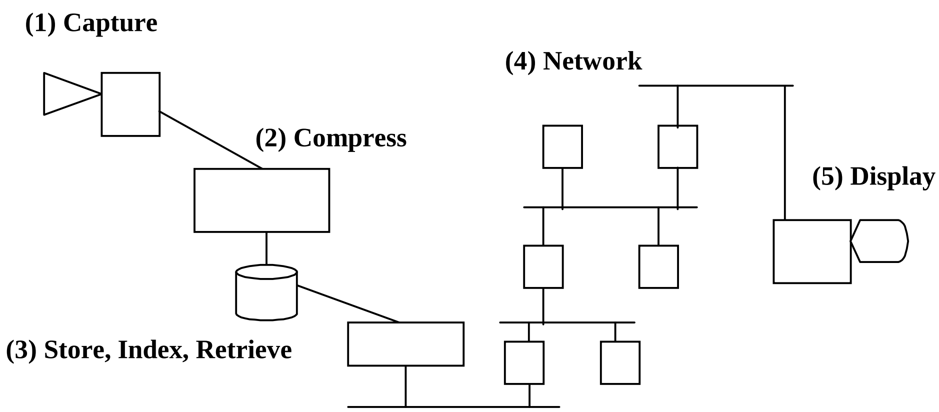

Figure 2: Quality of Service Multimedia Infrastructure

The equipment budget broken down by area is shown in Table 1 below.

Equipment Category	Five-Year Budget
Capture	$41,950
Compression	$102,000
Storage	$363,600
Networking	$360,850
Presentation and Development	$97,500
Total	**$965,900.00**

Table 1: Summary of Equipment by Category

In addition to supporting experiments conducted within research laboratories on the Purdue campus, our direct connection to the Internet 2 backbone via vBNS and Abilene will allow us to use this equipment to conduct extensive experiments with our colleagues at other research universities.

Figure 3, below, shows the laboratories at Purdue affected by this proposal. While the multimedia database laboratory in Math 431 will act as the focal point of the project, equipment supporting the research activities of the project will be installed in or used by people in the ten other labs identified in the figure.

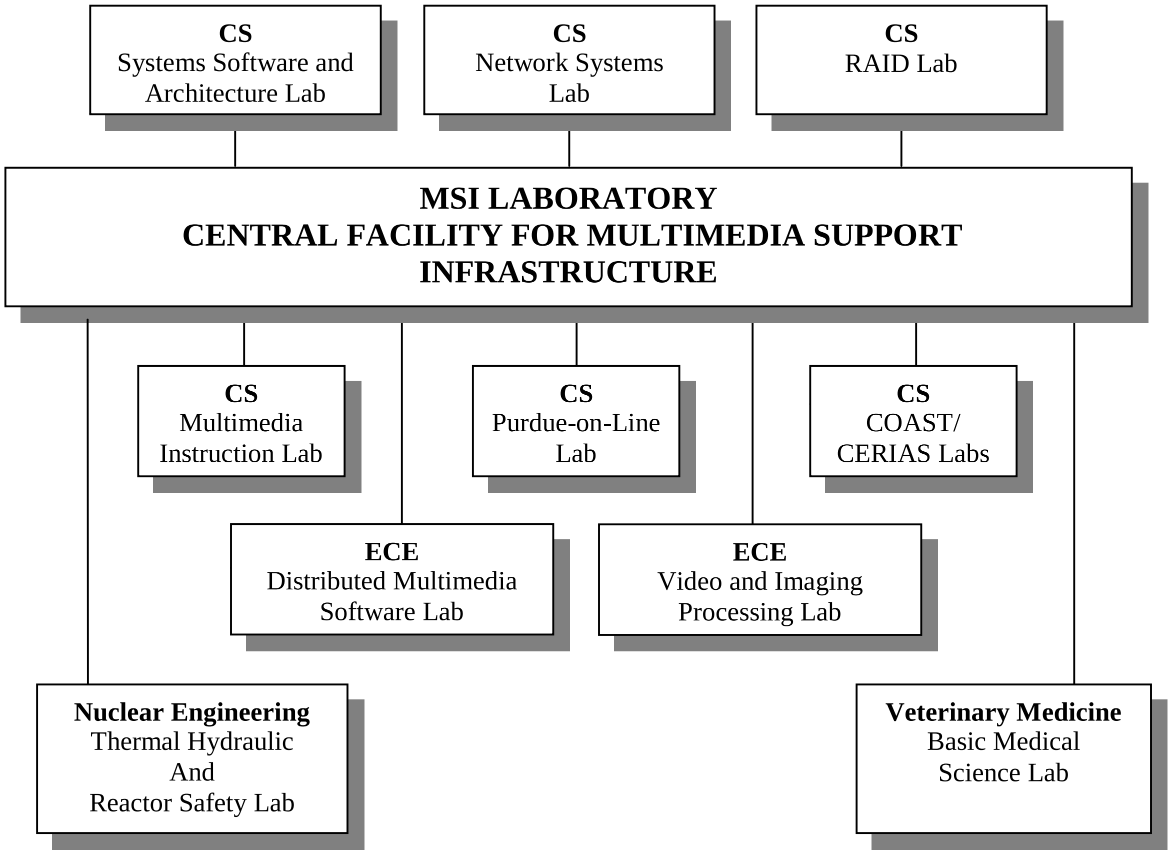

C.1 Networking

Delivery of multimedia data is time critical, requiring quality of service guarantees that depend on the capabilities of the underlying networks as well as the quality requirements of the delivered multimedia stream. High speed, lightly loaded networks permit the delivery of large image, full-motion video; low-speed or congested networks require scaling back of the video quality.

The proposed networking infrastructure consists of several high speed, state-of-the-art network packet transport technologies and advanced, programmable routers. The network technologies, including Gigabit Ethernet and ATM, will allow us to create a testbed that models the global Internet. The routers support quality of service protocols and techniques, such as weighted fair queuing and RSVP. Since they are programmable, we can use them to test our own modifications to the routing algorithms.

With this network equipment we will be able to build an Internet testbed, with traffic generators to simulate congestion conditions. Then, using our quality of service algorithms programmed into the routers, we can measure the expected throughput under typical Internet conditions. Since the routers are of the kind used on actual Internet backbones, we can also compare our algorithms to those used by the Internet under similar load conditions.

Qty	Description	Unit Cost	Total Cost
10	ATM switches	$7,000	$70,000
3	Gigabit Ethernet switch	$8,000	$24,000

1	Myrinet switch	$5,000	$5,000
4	Wireless network	$1,300	$5,200
10	WFQ and RSVP router	$21,665	$216,650
40	Special purpose network card	$1,000	$40,000
Total			**$360,850**

Table 2: Networking Equipment

C.2 Compression

Although we have researchers developing and testing new compression algorithms, we also want to use off-the-shelf compression hardware/software solutions to convert existing multimedia sources to digital form. These compression engines will allow us to not only take advantage of these standard technologies for normal use, but will also provide a benchmark against which to compare our new algorithms.

Qty	Description	Unit Cost	Total Cost
2	MPEG-1 encoder	$4,000	$8,000
12	MPEG-2 encoder	$7,000	$84,000
2	Digital video to MPEG encoder	$5,000	$10,000
Total			**$102,000**

Table 3: Compression Equipment

C.3 Storage Systems and Databases

We will use both commercial, off-the-shelf storage technologies and technologies developed by researchers in this project. The off-the-shelf technologies will provide a useful mass storage function for us during the initial years of the project. In later years, it will provide a benchmark against which to compare our own developing systems. The equipment request includes components to build and test storage systems.

We will also assemble the infrastructure necessary to build and experiment with our own high performance multimedia databases. This research includes techniques for indexing multimedia streams; we will support this research using a commercial database package.

Qty	Description	Unit Cost	Total Cost
2	Hierarchical storage management system	$59,000	$118,000
2	Database compute engine	$55,000	$110,000
2	RAID storage server	$17,000	$34,000
4	Experimental storage server	$25,400	$101,600
Total			**$363,600**

Table 4: Storage Systems and Databases Equipment

C.4 Capture and Presentation

We are not directly involved in research associated with the endpoints of the multimedia channel (video capture and presentation to the user), but we do need to both capture and present video to test and demonstrate our algorithms. Accordingly, we plan to acquire both consumer and industrial grade video cameras to create video streams, and video decoder hardware and software to display them. The cameras, both analog and digital, will be used especially to create video tapes of sources that pose interesting problems in video compression and delivery. These sources, such as those with rapid scene changes, present special challenges to video compression algorithms.

Qty	Description	Unit Cost	Total Cost
2	Analog video camera	$1,000	$2,000
10	Digital video camera	$1,300	$13,000
2	Digital video camera	$3,200	$6,400
2	Audio recorder	$1,000	$2,000
2	Still image recorder	$5,000	$10,000
2	Digital video player	$3,400	$6,800
25	Hardware MPEG-2 decoder	$700	$17,500
14	Tripod	$125	$1,750
10	Workstation	$3,000	$30,000
1	Server	$35,000	$35,000
5	Portable computer	$3,000	$15,000
Total			**$139,450**

Table 5: Capture and Presentation Equipment

D. Resource Allocation

D.1 Current Computational Facilities

Computer Science Department Facilities

The department is dedicated to providing high-quality computing facilities for use by computer science faculty, students, and administrative personnel. A technical staff operates the facilities and they are not only responsible for the installation and maintenance of the systems, but who also assist faculty and students in the development of software systems for research projects. The staff includes a director, facilities manager, administrative assistant, two hardware engineers, and five staff programmers.

General Facilities

General computing facilities are available for both administrative activities (such as the preparation of research reports and technical publications) and research needs that are not supported by other dedicated equipment. These facilities include a 32 processor SGI Origin 2000, eight Sun multiprocessor systems and several Sun and Windows NT file servers. The main systems each have 128 MB to 1 GB of main memory and a total of over 400 GB of disk storage. All faculty and many graduate students have a Sun, Silicon Graphics, Intel PC, or X display station on their desk.

Educational Facilities

Computing systems used by students enrolled in both undergraduate and graduate computer science courses include over 50 Intel Pentium-based PCs running either Sun Solaris or Windows NT. Two rooms in the Computer Science Building are dedicated to laboratory-based instruction using these facilities. A later section (Purdue University Computing Center Facilities) lists equipment owned and maintained by the Computing Center but used by computer science students.

I/O Equipment

The department operates both special-purpose output devices as well as general output equipment, including about 50 laser printers, color printers, color scanners, and video projectors.

Networking Services

The department is strongly committed to state-of-the-art networking technology to provide access to and communication among its systems, as well as to those elsewhere on campus and throughout the world. Twenty-five 10 Mb and 100Mb Ethernets included in the Computer Science Building connect the workstations and network terminal concentrators to the departmental computing facilities. An ATM fabric currently supports 45 workstations, and a fiber-optic ATM link connects departmental systems to other systems on campus, as well as to the Internet community. Experimental Gigabit Ethernets and Myrinets are also in use. ISDN and ADSL connections are in use for remote access from a number of nearby sites. Purdue University is a member of the Internet 2 project and has received NSF funding to connect to the vBNS/Abilene networks. With an Abilene gigaPOP located in Indianapolis, Purdue will have an excellent connection to Internet 2 and be well positioned to integrate its activities and experiments with this research network.

Purdue University Computing Center Facilities

In addition to the facilities described above, students and faculty have access to computing systems owned and operated by the Purdue University Computing Center (PUCC). General instructional facilities operated by PUCC include a large Sun SPARCservers, an IBM 3090, and dozens of Sun, Mac, and Windows laboratories. In addition, PUCC provides systems for use in courses taught by the CS Department. These systems include Sun SPARC 5 workstations for undergraduate computer science courses and Intel personal computers for use in introductory courses for non-majors.

Departmental research projects make use of other facilities provided by PUCC, including an Intel Paragon and IBM SP/2.

D.2 Description of Requested Equipment

Descriptions of the equipment, including the year in which each item is to be acquired is given in the table at the end of this section. In most cases, a sample vendor and model number of equipment matching the description is also included.

D.3 Rationale for Requested Equipment

Networking

The distribution of multimedia data will be done using high-speed networking hardware, including routers that support advanced quality-of-service protocols. Delivering simultaneous, independent streams of real-time video to many destinations requires not only high-speed networks, but also networks that use protocols that support differentiated services. This equipment will allow us to build the necessary infrastructure and experiment with new protocols and applications that use these protocols.

Acquisition of the networking hardware will be spread out over the lifetime of the grant. One of the workgroup switches will be used in a lab in Electrical and Computer Engineering and an other in a lab in Nuclear Engineering. These two switches will be connected to the experimental testbed in Computer Science using spare dark fiber currently available on campus.

Storage

We address two needs of the project with the purchase of (1) off-the-shell high-capacity storage hardware and (2) hardware components to allow us to build our own storage system. Capturing and compressing many multimedia sources will require us to have a large, commodity storage server. In addition, part of our research involves designing, building, and testing new large-scale storage systems to meet the real-time delivery requirements of streaming multimedia data. These purchases will allow us to store large data sets for indexing and analysis, and also to build novel, competing storage servers.

The storage equipment purchases will be spread out over three years. In the first year, we will buy a minimal hierarchical storage management system, augmenting it as necessary to keep up with our growing storage needs. In the second year, one of the database compute engines will be purchased for use in the School of Electrical and Computer Engineering. In the third year, a hierarchical storage management system will be added to that lab.

Compression

We will use off-the-shelf hardware to perform high quality, real-time conversion of video data to compressed form (primarily MPEG-2). We are requesting both “low-end” equipment for use in several research labs, as well as fewer “high-end” compression systems to produce high-resolution video streams. Some of the research Purdue will do includes the invention of competitive video compression techniques.

The purchase of the compression equipment will be evenly distributed over the first three years of the grant. One of the high end Optibase MPEG-2 encoders will be used in a laboratory in the School of Electrical and Computer Engineering.

Capture

The analog, digital, and still-frame recorders are needed to capture audio and video in classrooms, seminars, etc. This material will provide several sources of multimedia data with which to experiment. A high-end 3D still camera will be used by the School of Veterinary Medicine to capture images medical procedures.

Roughly ½ of the capture equipment will be purchased in the first year, with the balance purchased in the second year. This distribution will allow us to get started quickly building a library of multimedia data.

Presentation and Development

To present the video streams to the end-user applications requires video decoder hardware and software, and a small number of testbed workstations. The last step in video delivery needs special hardware and/or software to decode the incoming video stream in real time and display it to the user. This equipment will be needed to test the results of our research.

Much of the development workstations and decoder devices will be purchased in the first or second year of the grant. As the projects increase, more systems will be purchased to allow us to expand our experiments to more platforms.

E.3 Management Structure

We plan to establish three committees to oversee the equipment purchased in this proposal:

Steering Committee (project PIs): Insure that all project members have fair access to the equipment. Make final decisions on equipment purchases.

Advisory Committee (external researchers in same area): Make recommendations concerning research directions, equipment purchases, and infrastructure operation. We have obtained agreement from the following distinguished scientists to serve:

1. Avi Silberchatz, Director of the Information Sciences Research Center, Bell Labs.
2. Richard DeMillo, V.P. Of Research and General Manager of Information and Computer Science Research, Bellcore.
3. George Vanecek, Chief Scientist, AT&T Geoplex Systems.

Policy Committee (selected from project participants): Create access and use policies for approval by the Steering Committee.

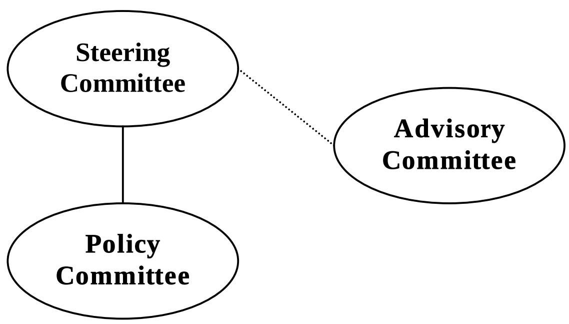

Figure 4: Management Committee Relationships

The intent is that the Steering Committee will be in charge of and responsible for overall activities and research on the grant. The Advisory Committee and Policy Committee will provide input to the Steering Committee. The Policy Committee will insure that all researchers, both those directly associated with the project as well as those collaborating with project members, are able to obtain fair access to the equipment.

Additionally, the committees will provide advice and direction to the departmental computing facilities staff, whose extensive experience and technical expertise in operating large collections of research equipment makes them ideally suited to install and maintain the proposed equipment. The director of the computing facilities for the Computer Science Department (Korb) will serve as the technical lead for the installation and operation of the project equipment. He will finalize the selection and integration of all project equipment. The director will also oversee the technical staff member hired by the project, as well as the departmental computing staff who will be responsible for the equipment.