SSCL-Preprint-66

SSCL-Preprint-66

Superconducting Super Collider Laboratory

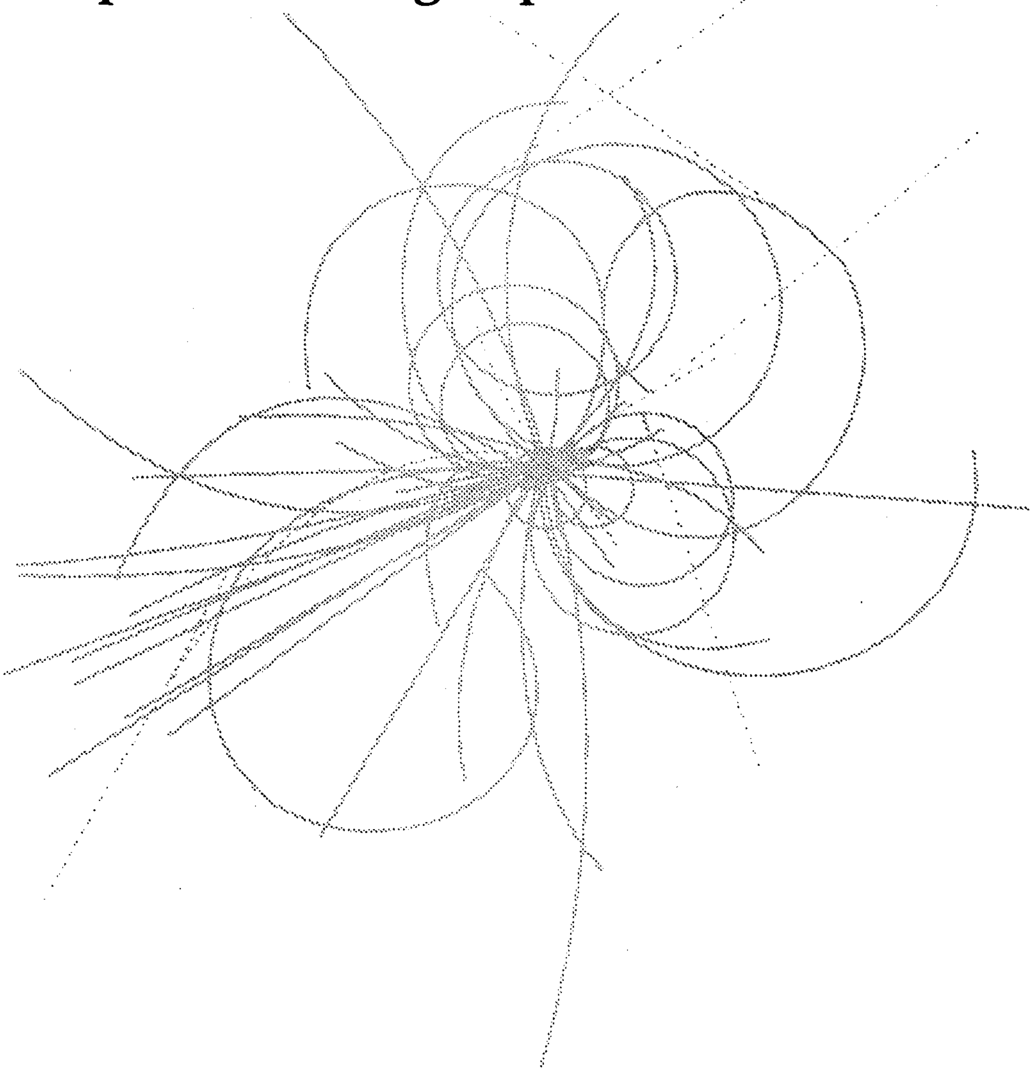

System Design Overview for the Magnet Test Database System

J. Ivey, X. Li, J. Garbarini, and P. Ball

March 1992

To be published in *Supercollider 4* SSCL-Preprint-66

System Design Overview for the Magnet Test Database System*

J. Ivey, X. Li, J. Garbarini, and P. Ball

Superconducting Super Collider Laboratory†
2550 Beckleymeade Ave.
Dallas, TX 75237

March 1992

*Presented at the International Industrial Symposium on the Super Collider, New Orleans, March 4–6, 1992.

†Operated by the Universities Research Association, Inc., for the U.S. Department of Energy under Contract No. DE-AC35-89ER40486.

SYSTEM DESIGN OVERVIEW FOR THE MAGNET TEST DATABASE SYSTEM

James Ivey, Xiaoyu Li, Joe Garbarini, Penny Ball

Superconducting Super Collider Laboratory*
2550 Beckleymeade Avenue
Dallas, TX 75237-3946

ABSTRACT

The Magnet Test & Data Management Department is responsible for testing of the superconducting magnets to be produced by the SSCL and industry, for data analysis of test results, and for creating and maintaining a central database of Laboratory magnet information. The magnet database will be accessed by magnet scientists at the SSCL, other HEP laboratories, and industry. The department uses a distributed and networked collection of UNIX workstations and servers to carry out its mission.

We will present an overview of the UNIX network, along with a history of design decisions and future plans for the system.

INTRODUCTION

Throughout the production phase in the MSD, the Test Department will be acquiring a large amount of data for later perusal and analysis. The Data Management Group of the MSD Test Department is developing and operating a data management system to integrate all this data. In the system support section, our responsibility is to provide a set of useful tools to aid in the interaction between the users and the database. At the same time however, it is necessary to keep our system as flexible and adaptable as possible, so as not to restrict the pursuits of the users. In fact, to the end-user our framework around the database should be so familiar and intuitive as to be unnoticed in everyday use. In our approach to this problem, we have focused on what we consider the four most important aspects of the system:

1) Computing and Network Resources
2) Local and Remote Access
3) Data Storage and Retrieval Options
4) Additional Safeguards and Security Measures.

*Operated by the Universities Research Association, Inc., for the U.S. Department of Energy, under Contract No. DE-AC35-89ER40486.

The purpose of this paper is to brief the reader on how we came to our decisions on the topics, and how we are attempting to implement those decisions. While a detailed description will not be possible in this short paper, reference documents are available from the MSD Test Department archive on a number of the topics[1].

SYSTEM OVERVIEW

Computing and Network Resources

The applications involved in this system will have very different and possibly very extensive needs. At the SSCL, a distributed processing environment has been a common solution to this problem[2]. To ensure that the requirements of each will not conflict with the other, or together ruin the performance of the entire system, it is necessary to divide these processing needs into discrete tasks and assign sufficient computing/network power to handle each task on a regular basis. We have decided to make this division of resources between the database services and the application programs which use that data. Our servers will be assigned the task of managing the database and providing file service to the rest of the network, while the individual workstations will be required to handle all user interaction and data analysis. This ensures that applications will not be in contention for processing time with the database system. As a side effect, it also allows us to select software and hardware for our applications without being restricted by the requirements of the database system, and vice-versa.

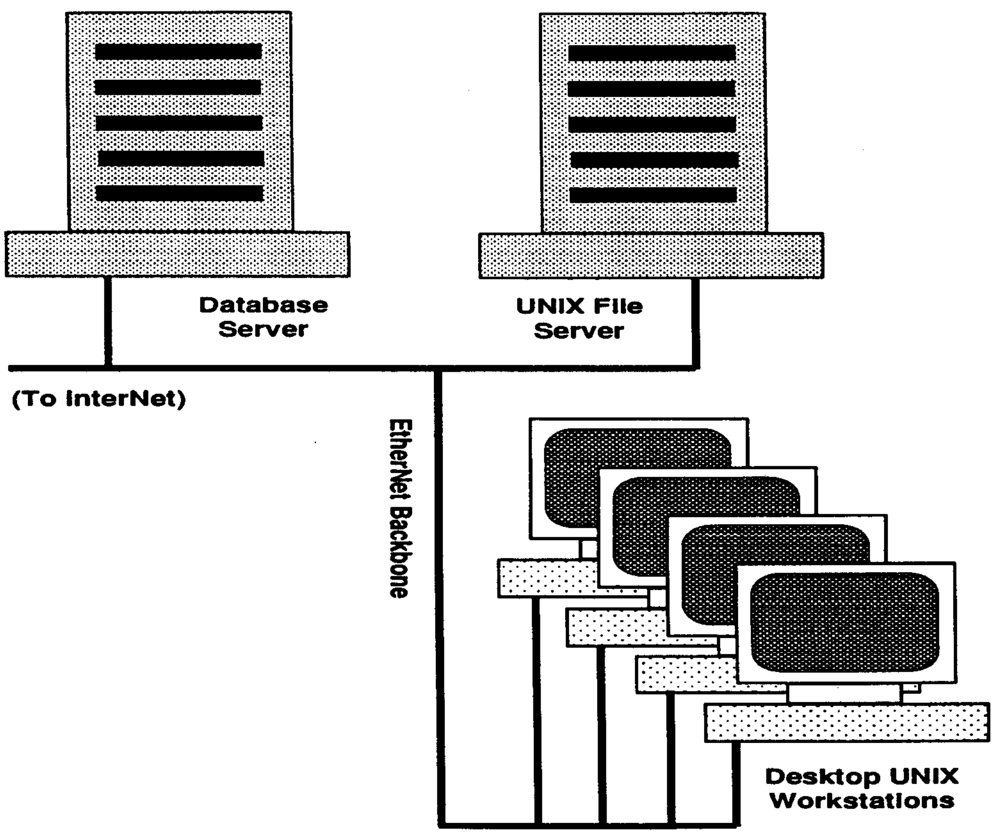

Figure 1. Current Layout of Network

Another decision, made to gain flexibility in our choice of application hardware and software, was to adopt UNIX as our primary operating system and X as the basis for selection of graphical user interfaces. While both have been criticized for various philosophical reasons, a greater segment of the commercial and scientific community is adopting them for their own uses. As such, we view UNIX and X as *adhoc* standards that give us a broad spectrum of choices when we look for products to fit our needs. X in particular has influenced a number of decisions in our ongoing design of the interfaces used to access our system.

Local and Remote Access

By moving the application processing away from the database server, we have given the users the full extent of our workstation network's processing power. In the past this service would have been available only to users on-site. However, our choice of X enables us to go a step further and provide the same functionality to users at remote locations. X allows a user to run an application on one computer with the user interface on a completely different system. To achieve this level of transparency to the user, X communicates using a protocol that handles the graphical events necessary for displaying the information, instead of the information itself. In this way, the excessive traffic that would be necessary to communicate our database information to other sites is removed, and is replaced with a subset small enough to present a fast interface without a dependency on physical location. Furthermore, because the underlying X protocol has been implemented by a large number of vendors, the interface can be displayed on very different architectures with little or no visible change.

One side effect of this decision has turned out to be one of the greatest benefits of the system. Because of the commonality of X implementations between various architectures, we are able to give all of our users a consistent, familiar interface. Our applications can be designed with a single GUI (Graphical User Interface), instead of customized for each specific environment. This is a tremendous help both to software engineers, who can focus on more critical design issues, and to end-users, who are not forced to relearn the interface when they move to a different computer or lab.

Data Storage and Retrieval Options

Although the choice to create this centralized database has a number of advantages that make it a vital addition to the SSCL,[3] it also has a number of inherent problems which we will need to address.

Size. When fully implemented, the size of our database, estimated at around half a terabyte (500,000 megabytes) of information, becomes a very restrictive criterion for selecting storage systems. While a number of technologies can handle this amount of data, each has its particular drawbacks that become more apparent when we examine our other storage problems. To interact fully with the database management system (DBMS), it is desirable to provide a storage system that will appear as a single entity to the software. All selection of files from the media must be transparent to the DBMS, allowing the greatest flexibility in our development of data retrieval and manipulation routines. No currently available storage option can handle our data requirements without the use of a collection of similar devices, known either as an array or a jukebox. Such devices must also provide a transparent software interface to integrate and manage the collection of media.

Lifetime/Reliability. Once of the strongest arguments for optical storage has been the extended shelf-life and reliability of its media. It is necessary to ensure that the data will still be intact at the end of the SSCL's 30-year lifespan. Magnetic storage is not currently capable of meeting this requirement without renewing the media every few years.

With both alternatives it is still necessary to investigate methods for back up. Once again optical seems to have an edge, as we would be able to copy the data from one set of optical media to another for back up purposes. With magnetic drives, however, we must find a separate system for back up, using magnetic tape or, possibly, a small optical WORM system.

Access. Because of the nature of the magnet test database and the uses it will be put to, it is necessary to investigate methods of making all the data available at once. This leaves us with two options: 1) Use a set of optical drives (known as a jukebox) to hold the raw data and keep the database structures and header information on magnetic storage, or 2) Use a set of magnetic drives linked together (known as a disk array) to hold all the information. The advantages and drawbacks were fairly clear. An optical jukebox, while the most cost-effective solution with the longest media shelflife, can be intolerably slow when switching between media. Also, while the technology has been around for a while, a standard format has not appeared in the marketplace, leaving us with doubts about the future of any system we purchase. A bank of magnetic drives, on the other hand, would provide us with the speed and transparent functionality that are a necessity for our project. All drives are on-line at the same time (in contrast to the optical storage, which requires switching between media at times) and may be accessed with no intermediary layer of software - only the normal UNIX I/O commands (another failing of most optical systems). In fact, the only drawbacks we see in a disk array are its cost and its need for a separate back up system. At the present time, we are still investigating both of these alternatives, though we are leaning toward a solution using a disk array.

Additional Safeguards and Security Measures

A final concern that has arisen during the initial implementation is how to ensure the integrity of our system, whether against accidental damage or from intruders with malicious intents. Unfortunately, because of the open philosophy of this system, we cannot simply close our doors to the outside. We must provide service to a very large user base, while somehow protecting against trespass. Additionally, we cannot restrict the users' environment without directly impacting their work.

Currently we are still providing a very open system. An account is available with a publicly available name and password. While the functionality of this account is restricted in a number of ways, the user still has the ability to access the database and generate queries and reports. Additionally, should a user need further access to the system, he may arrange for a personal account. These are supplied on a one-person/one-account basis, in order to maintain a line of responsibility for the account.

Over the next year we plan to obtain another server and our mass-storage solution. At this time we will retire one of our older servers to a role as the remote account server and mail host. In this way, it will become a "firewall" against the outside, providing a level of security that can be broken without any adverse effects on the rest of the network. While it will have the ability to read the database and other system files, it will have no ability to move to other systems or to write to the database. In all other respects, however, users will have accounts with the full functionality of a normal UNIX account. We think this will meet our modest security needs without hampering users.

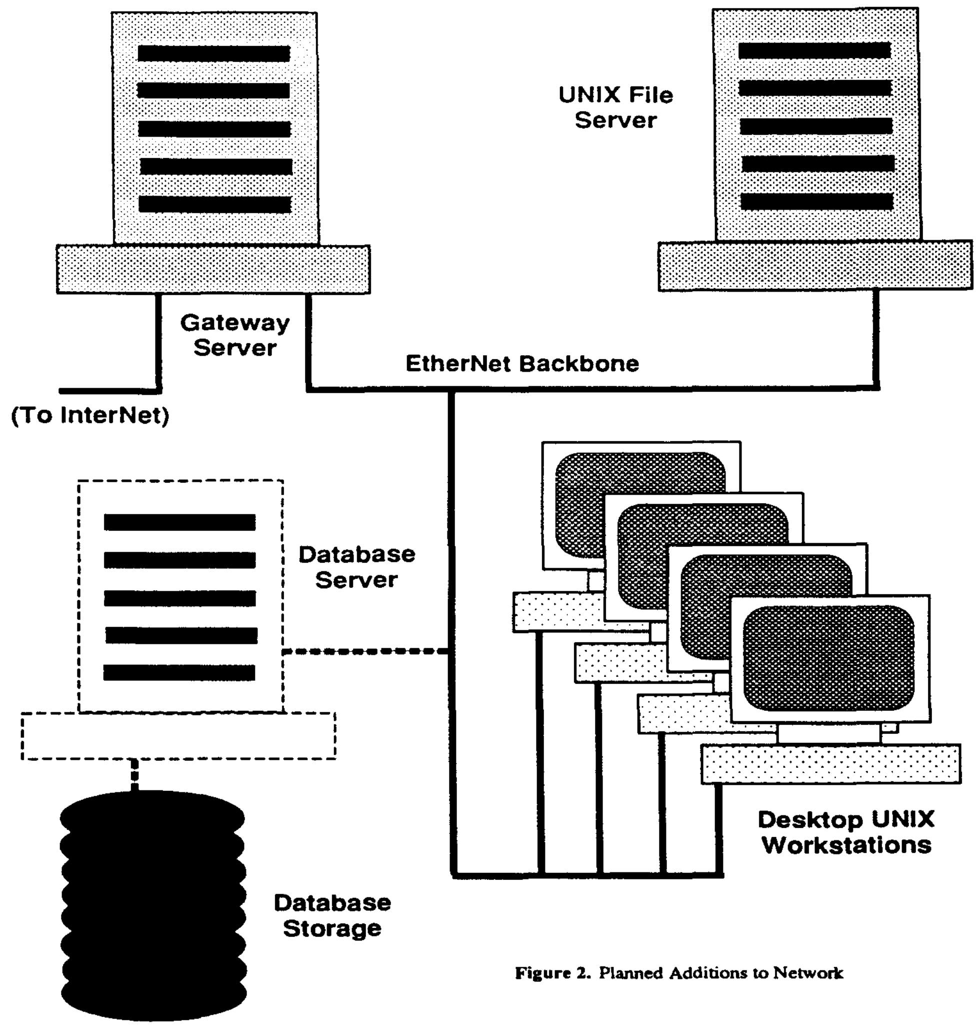

Figure 2. Planned Additions to Network

FUTURE PLANS/CONCLUSION

Unfortunately, after all our consideration we may still be forced to come to some compromise because of time constraints. Over the next two years we can expect 40-50 gigabytes from the magnets produced. With a number of computer vendors asking us to wait just a little longer for their new announcements, we have reached a point where we cannot afford to stay our decision anymore. Additionally, because of budget constraints, any system we obtain now must either be expandable to or already have the full capacity we will eventually need.

In response to the changing and unreliable marketplace, we have to accept a product's status as a standard as our primary criterion. In light of this, it is almost certain that we will pursue some sort of magnetic disk array system as our primary database storage, with the possible addition in the future of an optical WORM or magnetic tape back up system. In any event, it seems clear that some combination of the two technologies is necessary to meet the requirements we have laid out.

Our networking situation shows more promise, with a general acceptance by vendors of a standardization of network protocols and graphical user interfaces. In addition, with SUN Microsystem (our most common system) migrating to its new operating system, a System V-based UNIX, we find ourselves rapidly approaching a common environment on a number of levels. With careful planning in the next two years, we should find ourselves in an optimal marketplace for fully implementing the computing environment we have planned.

REFERENCES

1. Magnet Systems Division Test Department, MTL Documents 004-005, SSC Laboratory.
2. SSC Laboratory, "Information Resources Management FY 1994/1998 Long-Range Site Plan."
3. MSD Test and Data Management Group, "The Magnet Database System."