

National HPCC Software Exchange *

Shirley Browne, Jack Dongarra, Stan Green,
Keith Moore, Tom Rowan, and Reed Wade
University of Tennessee

Geoffrey Fox and Ken Hawick
Syracuse University

Ken Kennedy
Rice University

Jim Pool
California Institute of Technology

Rick Stevens, Bob Olsen, and Terry Disz
Argonne National Laboratory

January 9, 1995

Abstract

This report describes an effort to construct a National HPCC Software Exchange (NHSE). This system shows how the evolving National Information Infrastructure (NII) can be used to facilitate sharing of software and information among members of the High Performance Computing and Communications (HPCC) community. To access the system use the URL: <http://www.netlib.org/nse/>.

*Work on the National HPCC Software Exchange is supported by NASA under Grant No. NAG 5-2736. The NHSE is under development at Argonne National Laboratory, California Institute of Technology, Rice University, Syracuse University, and the University of Tennessee.

1 Introduction

The National HPCC Software Exchange (NHSE) is an Internet-accessible resource that will facilitate the exchange of software and information among research and computational scientists involved with High Performance Computing and Communications (HPCC). The purpose of the NHSE is to facilitate the development of discipline-oriented software and document repositories and to promote contributions to and use of such repositories by Grand Challenge teams, as well as other members of the high performance computing community.

The target audiences for the NHSE include HPCC application and computer scientists, users of government and academic supercomputer centers, and potential industrial users. The expected benefits from successful deployment of the NHSE include the following:

1. Faster development of better-quality software so that scientists can spend less time writing and debugging programs and more time on research problems.
2. Reduction of duplication of software development effort by sharing of software modules.
3. Reduction of time and effort spent in locating relevant software and information through the use of appropriate indexing and search mechanisms and domain-specific expert help systems.
4. Reduction of the time scientists spend dealing with information overload through the use of filters and automatic search mechanisms.

2 Overall Strategy for the NHSE

The effectiveness of the NHSE will depend on discipline-oriented groups and Grand Challenge teams having ownership of the discipline-oriented software repositories. The information and software residing in these repositories will be best maintained and kept up-to-date by the individual disciplines, rather than by centralized administration. Central administration will be used instead to handle interoperation and to meet common needs, such as indexing, searching, etc.

Although the various disciplines will have ownership of the repositories, they should not be expected to develop the software and tools for building, managing, and interfacing to their repositories. Because much information retrieval (IR) software is currently available, both in the form of client and server programs (e.g., http servers and WWW browsers), as much of this software should be incorporated into the NHSE as is useful. Using existing software has the following advantages:

1. It will reduce the amount of software development required for deployment of the NHSE.
2. It will ease interoperation with outside information sources that already use existing IR software.
3. It will take advantage of the familiarity that many users already have with these tools.
4. The NHSE will be able to take advantage of ongoing development efforts and enhancements for these tools without investing additional IR software development efforts of its own.
5. The HPC software and other information collected by the NHSE will be readily accessible by outside organizations such as schools, libraries, and museums.

The scope of the NHSE is software and software-related artifacts produced by and for the HPC Program. Software-related artifacts include algorithms, specifications, designs, and software documentation. A software package or software-related artifact is also called an *asset*, defined to be any item of interest stored in a reuse library. The following three types of software are to be made available:

- Systems software and software tools. This category includes parallel processing tools such as parallel compilers, message-passing communication subsystems, and parallel monitors and debuggers.
- Basic building blocks for accomplishing common computational and communication tasks. These building blocks will be of high quality and transportable across platforms. Building blocks are meant to be used by Grand Challenge teams and other researchers in implementing programs to solve computational problems. Use of high-quality transportable components will speed implementation, as well as increase the reliability of computed results.
- Research codes that have been developed to solve difficult computational problems. Many of these codes will have been developed to solve specific problems and thus will not be reusable as is. Rather, they will serve as proofs of concept and as models for developing general-purpose reusable software for solving broader classes of problems. The development of this reusable software is expected to be undertaken by commercial companies, rather than by academic researchers.

Although the different disciplines will maintain their own software repositories, users should not need to access each of these repositories separately.

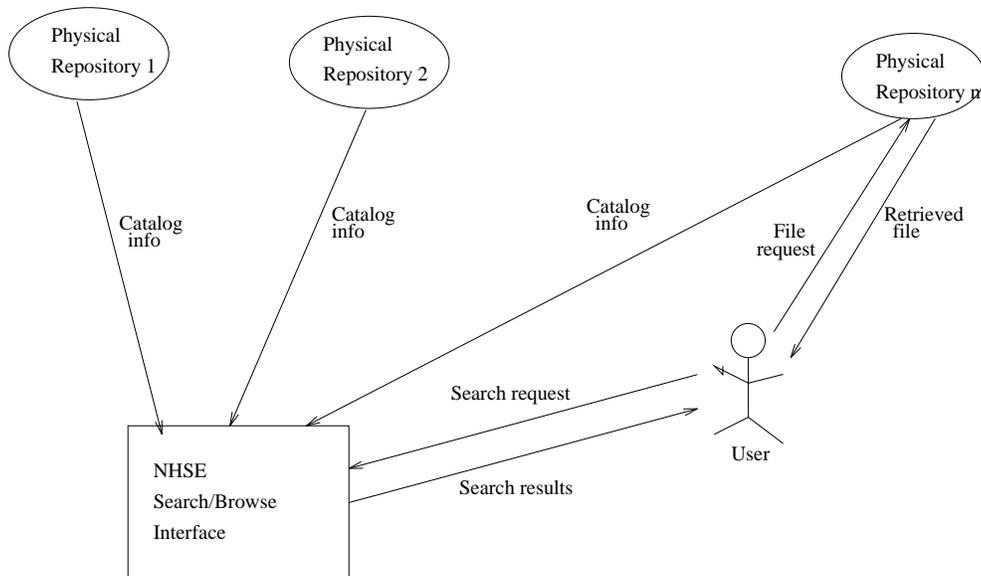


Figure 1: Virtual Repository Architecture

Rather, the NHSE will provide a uniform interface to a virtual HPC software repository which will be built on top of the distributed set of discipline-oriented repositories, as shown in Figure 1. The interface will assist the user in locating relevant resources and in retrieving these resources. A combined browse/search interface will allow the user to explore the various HPC areas and become familiar with the available resources. A longer term goal of the NHSE is to provide users with domain-specific expert help in locating and understanding relevant resources.

3 Audience

The intended audience for the NHSE consists of three communities:

1. The HPC application and computer science community.

These groups will be major sources of material for the NHSE. Members of the HPC applications community often need to develop highly optimized components from scratch, and thus generic reusable components may not be as useful to them as to other NHSE users. Higher level artifacts such as design documents, algorithms, and templates may be of use, however. Furthermore, the technique of domain analysis may be used to identify commonalities between applications, and the information thus uncovered can be used to assist the creation of new software systems.

2. Users of NASA, NSF, DOE, and other federal, state, and non-profit supercomputing centers.

These users may be involved in research, teaching, or industry. They will find libraries of reusable components valuable for providing good, if not optimal, solutions to common computational problems on high performance computing systems. They may also find tools such as parallelizing compilers and code restructurers useful for porting already-developed code to high performance computer systems. Educational users will find a wealth of materials that may be used for course enrichment and project ideas.

3. Current and potential industrial users interested in carrying out technology transfer.

The NHSE will promote software capitalization by providing access to extensive documentation on research prototypes developed in the academic community. The InfoMall USA technology transfer program ¹ developed by the Northeast Parallel Architectures Center at Syracuse University is an example of how such information may be provided and of how partnerships with industry may be formed [3].

4 Research and Development to be Carried Out

4.1 User Interface for Browsing and Searching

Current Web interfaces are difficult and frustrating for the user who is attempting to locate specific information. Browsing by following hypertext links is slow and can be disorienting. Keyword searching suffers from the vocabulary mismatch problem and is unsuitable for users with imprecise information and software needs. NHSE developers are working on support for an oriented, iterative combined searching and browsing process during which a user can

- augment her mental perception and knowledge of specific domains,
- improve her understanding of the problem she is attempting to solve, and
- form successively more focused and well articulated queries.

The interface will be in the form of a thesaurus-based roadmap. The NHSE will define the top levels of an HPCC thesaurus, drawing on an existing HPCC glossary [4] and on the current NHSE contents to generate thesaurus terms. Subject area specialists will be called upon to refine the lower levels. The thesaurus, along with a high-level classification scheme, will form the basis of a hypertext roadmap [6]. The roadmap will include scope notes and annotations to familiarize users with various HPCC areas and will serve as a springboard

¹ Accessible from a WWW browser at <http://www.npac.syr.edu/infomall/index.html>

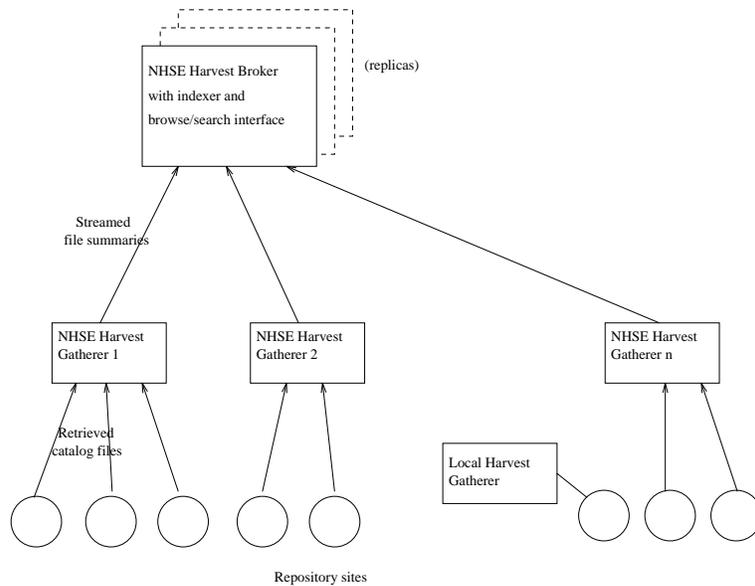


Figure 2: Indexing using Harvest

for thesaurus-assisted searches. Whether the HPCC thesaurus will function as both a searching and an indexing thesaurus, or only as a searching thesaurus, depends on the size and rate of growth of the HPCC software base, as well as on the available manpower for indexing, factors which have not yet been determined.

To enable searching, cataloging information must be made available for NHSE assets. Each physical repository will be responsible for maintaining one or more network-accessible file containing such cataloging information. These files will be retrieved and indexed by an NHSE indexer on a regular basis, and the resulting searchable index will be replicated for reliability. The NHSE will use the Harvest system [1] to do the collection, indexing, and index replication, as shown in Figure 2. Harvest components include the Gatherer, the Broker, and the Replicator. The Gatherer component retrieves and summarizes files containing cataloging information. The Broker periodically collects this information from the Gatherers, using an efficient stream protocol, and constructs a searchable index. A Gatherer can access an information provider's files across the network using the FTP, gopher, or HTTP protocol, or a Gatherer can be run locally by a provider site. The Harvest Replicator module can be used to replicate the Broker's index.

5 Naming and Authentication

Virtual access to a set of distributed, autonomously maintained repositories has many advantages but also poses numerous challenges. The main advantage of distributing the repository is to allow the software to be maintained by those in the best position to keep it up-to-date. Also, copies of popular software packages may be transparently mirrored to increase availability, improve response time, and prevent bottlenecks.

Many issues that are addressed in a centrally maintained repository by administrative procedures must be addressed in a virtual repository by other means. Such issues include assignment of unique identifiers to retrievable assets, collection and merging of cataloging information, and verification of the authenticity and integrity of retrieved assets.

The NHSE's approach to these issues will be to implement a location-independent naming architecture that unambiguously associates a unique name, called a Location Independent Filename (LIFN), with the byte contents of a published asset and that includes mechanisms for authentication and integrity checking. Authentication will ensure that the purported author of a published asset is the actual author; integrity checking will ensure that the contents of a retrieved asset are exactly the same as those published under the asset's unique name. Higher level names, called Uniform Resource Names (URNs), that are not associated with the specific byte contents, may also be assigned to assets.

Publishing tools will be made available to assist publishers with naming and cryptographically signing published assets, and with exporting asset descriptions to an NHSE search service. A distributed name-to-location lookup service will be provided, along with a means for publisher and mirror sites to register locations for published assets. A client library that may be linked with a WWW browser to enable the browser to resolve location-independent names and to perform authenticity and integrity checking will also be provided.

The steps involved in publishing an asset are shown as P1, P2, and P3 in Figure 3. A mirror site that maintains an authorized copy of an asset may also register a location for that asset. The steps a user carries out in searching for and retrieving a published asset are shown as U1, U2, U3, U4 in Figure 3. Having unique verifiable names will allow search services to unambiguously associate descriptions, including third-party descriptions such as critical reviews, with published assets. Scientific researchers will be able to unambiguously refer to software used to produce experimental results. Different users who have downloaded copies of a software asset that have the same name, or a user who downloads the same named asset more than once, will have the assurance that the copies are indeed the same. Unique naming will also facilitate collection management and tracking of assets by file servers and search services.

Authentication of assets will be handled by an asymmetric public-private key encryption system. A publisher will sign his description of an asset using his private key. Then any client program in possession of the publisher's public

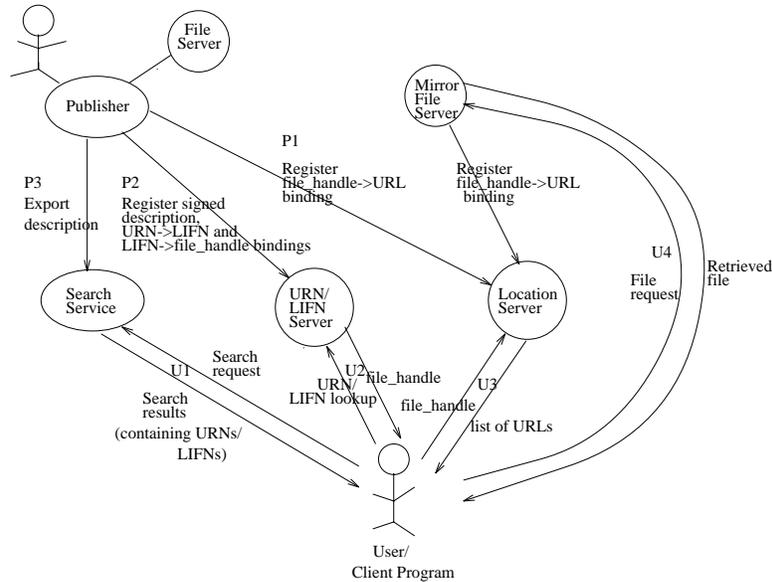


Figure 3: Publishing and Retrieving Assets Using Location-Independent Naming

key will be able to authenticate the asset description. Either the name or the description for an asset will include a signature for the file containing that asset, such as the MD5 fingerprint. A client program will be able to perform an integrity check on a retrieved file by computing the signature for the file and comparing it with the one known to be associated with that asset. These mechanisms are similar to those proposed in [5].

To avoid the overhead of the client having to perform authenticity and integrity checks for every file accessed, the NHSE plans to use an authentication system for name-to-location servers and for file servers. Name-to-location servers will allow only authorized trusted file servers to register locations of files. A trusted file server will guarantee that the file it returns for a particular name is correct. Updates from file servers to name-to-location servers, as well as the update protocol between replicated name-to-location servers, will require authentication. Such authentication may be based on public keys, shared secrets, network addresses, or some combination of these. Thus, authentication will be provided at the server level, rather than only at the individual file level.

5.1 Quality Control

Users of the NHSE need to have confidence that the software they obtain is of high-quality and well-tested. If the software is experimental or untested, they

should be made aware of this. Quality control will be impossible to automate completely because it requires human judgement and evaluation. Tools and procedures need to be developed that facilitate the testing and classification or labeling of software with respect to its quality and reliability. Research is needed to determine what quality control information is most useful and can reasonably be obtained, how to acquire this information, and in what format it should be stored for easy access.

5.2 Interoperation

The NHSE will catalog and provide access to software and software-related artifacts from all the HPCC software repositories. Assets accessible from other existing software repositories, such as ASSET, CARDS, DSRS, and ESLA, to name just a few, may also be of interest to NHSE users. Likewise, users of these other repositories may be interested in NHSE software. The NHSE will be participating in a small-scale interoperability experiment with the above repositories that will help define requirements for further interoperation efforts. The NHSE will also be working with the Reuse Library Interoperability Group (RIG) on establishing standards for unique naming, asset description and classification, and asset evaluation.

6 Current Status

An Internet-accessible prototype of the National HPCC Software Exchange has been developed ². The goals of the prototype are the following:

- Show the feasibility of distributing software and information via the National Information Infrastructure.
- Demonstrate access to existing software repositories.
- Demonstrate browsing and searching mechanisms.
- Identify limitations of current information retrieval techniques and areas needing further development.

The prototype NHSE was designed to be accessed from an forms-capable World Wide Web (WWW) browser, such as recent versions of NCSA Mosaic or Netscape Communication. The information and software were either already available or were made available using existing information access systems, such as FTP, Gopher, WAIS, and HTTP.

A user-friendly interface has been constructed in the form of a distributed HTML document. The root of this document is the NHSE home page, which

²Accessible from a WWW browser at <http://www.netlib.org/nse/>

contains an outline from which users can reach all the collected HPC information with a few clicks of the mouse. Users may also perform a keyword search on a free-text index of this distributed document.

The amount and diversity of material already available are too great to be able to provide a comprehensive summary here. Rather, descriptions of a small sample will be given to provide you with some idea of what is available.

The NHSE includes Netlib³, which is a moderated collection of high-quality mathematical software and parallel programming tools. Netlib contains some packages specifically designed for high performance computing, such as LAPACK and ScaLAPACK, which are libraries of numerical linear algebra sub-routines [2]. LAPACK is designed to run efficiently on shared-memory vector and parallel processors, while ScaLAPACK is designed for distributed-memory multiprocessors. Netlib also contains a number of parallel programming environments, such as MPI, P4, PICL, PRESTO, and PVM.

The NHSE includes pointers to a number of other established software repositories, such as ASSET, CARDS, DSRS, and ELSA. In the future, the NHSE will interoperate with these other repositories so that software from them may be retrieved directly from the NHSE interface.

The NHSE also points to a number of independently-provided specialized software collections. For example, the software exchange site maintained by the HPC Group at the University of Washington Astronomy Department⁴ contains tools for visualizing and doing calculations on output from cosmological N-body simulations.

7 Need for User Input

To develop a maximally effective searching and browsing interface, we need help from our users. We are developing a thesaurus-based roadmap on which to base our browsing and navigation tools, but the thesaurus vocabulary should be derived from your query terminology, needs, and potential uses of HPC technologies. We need to be able to make the connection between the problem you need solved and the NHSE resources that will help you solve it. Whether you are a government, industry, or academic user, and whether you are involved in teaching, research, technology transfer, or other activity that would benefit from HPC technology, we need to hear from you. Please visit the NHSE home page⁵ and fill out a user profile. To the extent possible with our available manpower, we will attempt to send you a personalized reply informing you of resources we are currently aware of that may be of use to you, and we will keep you posted of future developments. You will also be contributing to the development of a valuable resource for you and other HPC users.

³ Accessible from <http://www.netlib.org/>

⁴ Accessible from <http://www-hpcc.astro.washington.edu/>

⁵ Accessible at <http://www.netlib.org/nse/>

References

- [1] C. M. Bowman, P. B. Danzig, D. R. Hardy, U. Manber, and M. F. Schwartz. Harvest: A scalable, customizable discovery and access system. Technical Report CU-CS-732-94, Department of Computer Science, University of Colorado-Boulder, Aug. 1994.
- [2] J. J. Dongarra and D. W. Walker. Software libraries for linear algebra computation on high-performance computers. Technical Report ORNL TM-12404, Oak Ridge National Laboratory, Aug. 1993.
- [3] G. C. Fox, W. Furmanski, K. A. Hawick, and D. M. Leskiw. Exploration of the InfoMall concept. Technical report, Northeast Parallel Architectures Center at Syracuse University, Aug. 1994.
- [4] K. A. Hawick. High Performance Computing and Communications glossary. Technical report, Northeast Parallel Architectures Center at Syracuse University, July 1994.
- [5] J. W. Moore. The use of encryption to ensure the integrity of reusable software components. In *Proc. Third International Conference on Software Reusability*. IEEE Computer Society Press, Nov. 1994.
- [6] R. Pollard. A hypertext-based thesaurus as a subject browsing aid for bibliographic databases. *Information and Management*, 29(3):345–357, 1993.