Y/AMT-390

Y-12

OAK RIDGE
Y-12
PLANT

LOCKHEED MARTIN

Project Accomplishment Summary
for
Project Number 93-MULT-052-C1-04

DEVELOPMENT OF A HIGH PERFORMANCE STORAGE SYSTEM (HPSS)

Kenneth L. Kliewer
Lockheed Martin Energy Systems, Inc.
Dennis E. Stoker
IBM Government Systems

December 27, 1996

Approved for Public Release;
distribution is unlimited.

Prepared by the
Oak Ridge Y-12 Plant
managed by
LOCKHEED MARTIN ENERGY SYSTEMS, INC.
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-84OR21400

MANAGED BY
LOCKHEED MARTIN ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

UCN-13672 (26 6-95)

PROJECT ACCOMPLISHMENT SUMMARY

Title: Development of a High Performance Storage System (HPSS)
DOE TTI Number: 93-MULT-052-C1-04
CRADA Number: Y1293-0203
Partner: IBM Government Systems

BACKGROUND

An essential component in a high performance computing environment is a data storage and access system capable of high speed data transfers. Prior to the initiation of this project, all storage system software packages could, at a given time, move data only between two devices, for example, a computer and a tape drive or a disk and a tape drive. This limitation on overall system performance was unacceptable. The storage/access aspects of high performance computing input/output (I/O) had been badly neglected while individual Central Processing Unit (CPU) performance soared and the capabilities of parallel computing, with hundreds and even thousands of CPUs working together, was pushing computer performance to unprecedented levels. The need for high performance storage software capable of moving data across multiple paths simultaneously (i.e., parallel) was evident.

The software development backgrounds and capabilities of the major DOE computing laboratories (Los Alamos National Laboratory, Lawrence Livermore National Laboratory, Sandia National Laboratories, and LMES/ORNL at the beginning of the project, now LMER/ORNL) together with their need for high performance storage systems in their own computing environments compelled the initiation of this project. The IBM presence in the realm of data storage, both hardware and software, made clear that a partnership with them would be both effective and mutually beneficial.

DESCRIPTION

The overall objective of the project was the development of a parallel high performance storage software package capable of data transfer rates above 1 gigabyte/sec with files of essentially unlimited size. This necessitated modules for uniquely identifying files to be stored, for establishing the appropriate locale for the file in the storage hardware, for moving the file in parallel to the selected locale, and for making possible ready access to the file when desired. And all of this must be done with absolute accuracy and reliability while ensuring security at the requisite level. Responsibility for the various modules was distributed across the participating laboratories. The central LMER responsibility was the Storage System Management (SSM) package, the software package that controls all storage and access activities and provides readily understandable and complete information concerning system status to an operator. This information includes storage and access activity in progress; the location, size, and character of all files; and warning and error messages, among others. As such, SSM must be tightly coordinated with all of the HPSS modules and components and must represent, in effect, a synthesis of all. The result of this very extensive LMER effort was an SSM system that required approximately 83,000 physical lines of computer code.

As noted previously, this CRADA is a strongly coordinated project involving IBM Government Systems together with the aforementioned DOE computing laboratories. The LMER responsibility, the SSM, requires a synthesis of all the developmental projects. The IBM role in

the SSM was to coordinate the overall HPSS project and to ensure that all necessary concepts and features were brought to the attention of the SSM team and included in the SSM structure. This required extensive and continuous e-mail communications with all groups, comprehensive analyses, participation in lengthy tele-conferences of the HPSS Technical Committee, and participation in the wide ranging decision processes that the SSM required.

The product of this widely distributed, but remarkably well coordinated, project was Release 2 of HPSS, a storage system software package that met all of the initial goals.

BENEFITS TO DOE

As noted above, all of the laboratories participating in the development of HPSS required the capabilities that this new software system provided. This alone justifies the HPSS expenditures. However, there has emerged a benefit to DOE of overwhelming importance. The fact is that the only storage software package on the horizon that can meet the requirements of the ASCI project is HPSS. This makes the current availability of HPSS extraordinarily timely and provides a remarkable endorsement of the decision to initiate the HPSS project.

ECONOMIC IMPACT

HPSS is now announced as an IBM product and, accordingly, has the benefit of the IBM marketing organization. HPSS is in operation at the DOE laboratories noted, at the University of Washington, at the Fermi National Accelerator Laboratory at Cornell University, and at the Maui High Performance Computing Center, among others. Major computer vendors are exploring the use of HPSS in their platforms. Interest in HPSS from prominent high energy physics laboratories and weather stations world wide is extensive and growing. All indications are that the economic impact of HPSS will be substantial.

PROJECT STATUS

The HPSS project is continuing. The CRADA ended with Release 2. Release 3 appeared in the summer of 1996, with Release 4, now being defined, scheduled for the spring of 1997. Primary development support is now coming from the ASCI Project, and is extensive.

CONTACTS

Lockheed Martin Energy Research
K. L. Kliewer
Oak Ridge National Laboratory
P. O. Box 2008
Oak Ridge, TN 37831-6203

IBM Government Systems
Dennis E. Stoker
Phone: 713-335-4181
Fax: 713-335-4231

PROJECT EXAMPLES

The main product of this project is the huge HPSS code, which is proprietary. However, there are a number of manuals describing various aspects of operational procedures and capabilities

that are available. An HPSS tutorial is available on the World Wide Web at http://www.ccs.ornl.gov/hpss.

TECHNOLOGY COMMERCIALIZATION

As noted above, HPSS is now an officially recognized IBM product.