Available online at www.sciencedirect.com
SCIENCE DIRECT®
Fusion Engineering and Design 71 (2004) 89–93
Fusion Engineering and Design
www.elsevier.com/locate/fusengdes

Migration of alcator C-Mod computer infrastructure to Linux

T.W. Fredian*, M. Greenwald, J.A. Stillerman

Massachusetts Institute of Tech., MIT Plasma Fusion and Science Center, 175 Albany Street, Cambridge, MA 02139, USA

Available online 28 May 2004

Abstract

The Alcator C-Mod fusion experiment at MIT in Cambridge, Massachusetts has been operating for twelve years. The data handling for the experiment during most of this period was based on MDSplus running on a cluster of VAX and Alpha computers using the OpenVMS operating system. While the OpenVMS operating system provided a stable reliable platform, the support of the operating system and the software layered on the system has deteriorated in recent years. With the advent of extremely powerful low cost personal computers and the increasing popularity and robustness of the Linux operating system a decision was made to migrate the data handling systems for C-Mod to a collection of PC's running Linux. This paper will describe the new system configuration, the effort involved in the migration from OpenVMS, the results of the first run campaign under the new configuration and the impact the switch may have on the rest of the MDSplus community.

Keywords: Data acquisition; Computers; MDSplus; Diagnostics; Alcator C-Mod

1. Introduction

For the past 20 years, the computers used for fusion research at the Plasma Science and Fusion Center at MIT were VAX or Alpha systems running the OpenVMS operating system. The vendor for this hardware and operating system has undergone numerous restructuring over this time frame and the quality of the both the hardware and operating system has deteriorated. For this reason it was decided to migrate to another computing platform. The Linux operating system was chosen and this paper will discuss our experiences in the migration of the C-Mod data handling system from OpenVMS to Linux.

2. Background

The OpenVMS operating system has been used as the main computing platform for the fusion program at MIT for more than 20 years. Several major experiments including Alcator C, Tara, and Alcator C-Mod as well as several smaller programs used OpenVMS systems for data acquisition and analysis. During this time frame both the MDS [1] and MDSplus [2] data systems, widely used throughout the fusion research community, were developed under this operating system.

Changing hardware architectures along with corporate restructuring due to the acquisition of Digital Equipment by Compaq and Hewlett Packard has taken its toll on the viability of OpenVMS. Once a popular and robust computing platform, OpenVMS on VAX or Alpha computers has deteriorated in performance and

* Corresponding author. Tel.: +1-617-253-7623; fax: +1-617-253-0627.
E-mail address: twf@psfc.mit.edu (T.W. Fredian).

0920-3796/$ – see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.fusengdes.2004.04.016

capability relative to competing systems over the last 5–10 years and it became increasingly clear that we would need to move away from the OpenVMS operating system as the main computing infrastructure for the Alcator C-Mod experiment.

The first step was a major porting effort of MDSplus to UNIX and Windows platforms which was carried out in 1998–1999 by MIT, GA, LLNL, PPPL and Padova. Many sites have since adopted MDSplus for data handling or to provide remote access to their existing data stores using non-OpenVMS computing platforms [3]. With the port essentially complete and the availability of inexpensive but powerful personal computers, it was decided to experiment with Linux workstations for data analysis and visualization in the C-Mod control room. This experiment performed during the 2001 campaign was quite successful and led to the decision to migrate entirely from the OpenVMS based computing infrastructure of the Alcator C-Mod experiment to one based on the Linux operating system.

3. Migration—significant challenges

Switching away from OpenVMS at C-Mod was not going to be an easy task. Most of the C-Mod computing staff and many of the scientists and engineers had over 20 years of experience using the OpenVMS operating system while possessing little or no experience with Linux. The OpenVMS computing infrastructure was quite large with over 30 high-end physics and engineering workstations along with several server systems. Essentially all of the data acquisition software which communicates directly with the measurement devices such as CAMAC modules would need to be ported from OpenVMS to Linux.

In addition to the software libraries and applications requiring porting, a whole new set of system management procedures were needed. Procedures for account management, data archival, system backups, system configuration and replication, system monitoring and tuning would need to be developed.

Both users and the computer support staff would have to learn how to use and maintain a relatively unfamiliar computing infrastructure. There were many unknowns regarding robustness of the operating system as well as the software used on these systems, including MDSplus, when placed under heavy load during full scale data acquisition and analysis. There were also concerns about the ability to diagnose and repair problems should they arise during experimental operations.

4. Migration strategy

A migration strategy was formulated to minimize the risks and reduce inconvenience for the users. This staged migration plan began with the purchase of a few Linux workstations. Engineers and scientists could use these workstations to familiarize themselves with the Linux operating system and to port some of their analysis and visualization applications. This stage was successfully carried out toward the end of 2001. The next phase was to purchase more Linux based workstations and actually replace a third of the OpenVMS. A subset of the physics staff volunteered to use the Linux workstations for analysis and visualization during a C-Mod run campaign. This phase was completed mid-2002 and again was quite successful. The new workstations performed well, outperforming the more expensive but much older OpenVMS systems. Users were quickly gaining some expertise with Linux and found a wealth of information on the web to help them when they ran into difficulties. The ease of transition along with the increased performance of the workstations allowed us to take the next step in the migration, to replace most of the OpenVMS workstations and to move the data acquisition and data file management from OpenVMS to Linux servers. A new Linux server and another 12 Linux workstations were purchased shortly before then next C-Mod run campaign which began in March 2003.

Until this point the experimental data store was located on OpenVMS servers with the Linux workstations using the mdsip protocol over tcpip to access the data. The next step was to move the data store from the OpenVMS server to Linux. The new Linux server was configured with 1.5 terabytes of RAID storage and a copy of the entire C-Mod data archive was copied to it. Both the OpenVMS systems and the Linux workstations were configured to access C-Mod data via the Linux server. The original copy of the archive remained on the OpenVMS server in case there was a problem that would make it necessary to revert back to

the original configuration. This step was successfully completed toward the end of 2002 during an extended outage between experimental campaigns.

The next major step was to move the bulk of the data acquisition systems as well as the action dispatching and action servers from OpenVMS to Linux. This included getting the CAMAC serial highway software fully functional on Linux and the port of the numerous CAMAC module data acquisition support codes. The RAID storage on the OpenVMS server would need to be emptied and moved to the Linux server to provide sufficient storage for the upcoming C-Mod run. Switching the CAMAC hardware, the action dispatcher and action servers which perform data acquisition and analysis, and the location of the data to Linux were critical activities as it became increasingly difficult to revert to VMS in the event of some unforeseen difficulty. While a series of tests were carried out to evaluate the operation of the new configuration, it was unknown how well it would perform under full load during a C-Mod run campaign when the large number of simultaneous users puts significant stress on the overall system.

5. Results

The C-Mod began operation in March 2003 using a predominantly Linux based computing infrastructure consisting of 2 Linux server systems and 22 Linux physics and engineering workstations. Several software problems were discovered during the initial machine cycles when bringing the experiment online after an extended outage. Most of these problems were related to the action servers and dispatching which are multi-threaded and susceptible to race conditions and reentrancy problems. There were problems in some of the CAMAC support code as well but essentially all of the problems were fixed before full plasma operation of the C-Mod experiment began.

The relocation of the MDSplus data files from an OpenVMS server to a Linux server enabled us to use the “distributed” tree access mode when referencing data from the workstations. MDSplus supports three different types of remote data access called “thin client”, “thick client” and “distributed” and all three are based on the mdsip protocol which uses tcpip sockets for transport. Only the “thin client” and “thick client” access methods are supported with an OpenVMS server. In “thin client” mode the server performs all expression evaluation and data decompression. This mode places a high compute load on the server and also results in high network loads. In “thick client” mode, the server decompresses the data but the expression evaluation is done by the client. This reduces the computing load on the server somewhat. In “distributed” mode, the server only performs file I/O on behalf of the client. This results in the least compute load on the server and the smallest load on the network. It also support truly distributed trees where the data files can reside on numerous servers. While a similar functionality could be accomplished with file sharing protocols such as NFS or AFS, these solutions tend to be extremely inefficient when there are multiple concurrent readers and writers of the same data files which is a key feature of the MDSplus system.

The performance and robustness of the new configuration far exceeded expectations. The data acquisition rate was almost twice as high as our OpenVMS systems and the reliability of the data acquisition and analysis systems was as good as or better than what we experienced on OpenVMS.

The users of the system have been pleased with the increased performance in the analysis and visualization of data. Many of the analysis programs run in a fraction of the time that they used to run on the older OpenVMS machines. While the users have had some difficulties adjusting to the new computing environment, they also are pleased with the suite of applications available and the wealth of information on the Internet available to assist them in learning.

6. Final migration steps

The last run campaign of the Alcator C-Mod experiment ending in July 2003 was operated on a computing environment consisting predominantly of Linux systems. There are still a few applications which still run exclusively on OpenVMS which will be moved to Linux by the end of the 2003 calendar year. The most demanding of these applications are those associated with the real-time plasma control system currently used on C-Mod. The plasma control for C-Mod [4] is based on a hybrid analog/digital system developed by the EPFL in Lausanne. The host system (OpenVMS)

uses Bitbus to communicate with the hybrid control system. Work is underway to replace this entire system with a digital solution based on data acquisition cards running on CPCI (compact personal computer interface) under the control of a PC running Linux.

Assuming the final steps of migration go as smoothly as the migration completed thus far we anticipate retiring the remaining OpenVMS systems used by the C-Mod experiment by the end of 2003.

7. Conclusions

The OpenVMS operating system has provided a robust computing environment for the Alcator C-Mod experiment for more than a decade. Both the MDS and MDSplus systems used widely in the fusion research community were developed on the OpenVMS platform. While OpenVMS has served the community well for many years, advances in hardware and software as well as market forces have lead to the end of OpenVMS as a preferred platform. Linux, on the other hand, is gaining in acceptance in both the scientific and business environments. It runs on inexpensive yet increasingly powerful personal computer systems which now have the computing capacity rivaling or exceeding that of systems classified as super computers only a few years ago. For these reasons we decided to begin the transition from OpenVMS to Linux as a computing environment for the Alcator C-Mod experiment. As we began the transition it was assumed that we might end up with a combination of OpenVMS servers and Linux workstations but as the transition took place it became clear that the robustness and speed of Linux and PC's were sufficient to make the retention of OpenVMS unnecessary and undesirable.

We are well on our way toward a complete conversion to Linux as a computing platform for Alcator C-Mod. The migration thus far has been more successful, more rapid and more trouble free than we anticipated. Performance of the C-Mod data system has been more than adequate and has been able to keep up with the increasing data handling demands now exceeding 650 megabytes of data per machine cycle.

The migration from OpenVMS to Linux at Alcator C-Mod should have additional positive side effects. Most of the MDSplus core software is maintained at MIT and the migration should improve the quality of the MDSplus software on the Linux, other UNIX variants and Windows platforms which share the same code base. Several problems have been detected and fixed since the migration began as more of the code is being exercised. The experience we have gained in performing this migration can be used to assist other sites which may be performing similar migrations in the future [5,6].

At the same time, MIT is committed to continue support of MDSplus under OpenVMS and will retain systems to perform this task as necessary even though the OpenVMS version of MDSplus has been quite stable for several years.

Acknowledgements

Many people contributed to making this transition so successful. We would like to thank the MDSplus developers from GA, PPPL, LLNL and Padova [7] who assisted MIT in the port of MDSplus to UNIX and Windows. We would also like to thank all the MDSplus users and developers from numerous locations who bravely adopted MDSplus on computing platforms other than OpenVMS and served as the initial testers of the package on these platforms and provided invaluable information to assist in the debugging of the software. Last but not least we would like to thank the physicists, engineers and managers of the Alcator C-Mod experiment for their patience, encouragement and hard work during the transition period.

References

[1] T.W. Fredian, J.A. Stillerman, MDS/MIT high speed data acquisition and analysis system, Rev. Sci. Instrum. 57 (1986) 1907–1909.

[2] J.A. Stillerman, T.W. Fredian, K.A. Klare, G. Manduchi, MDSplus data acquisition system, Rev. Sci. Instrum. 68 (1997) 939–940.

[3] D.P. Schissel, et al., Enhanced computational infrastructure for data analysis at the DIII-D National Fusion Facility, Presented at the 2nd IAEA Technical Committee Meeting on Control, Data Acquisition, and Remote Participation on Fusion Research, Lisboa, Portugal, 19–21 July 1999, Fusion Eng. Design 48 (2000) 105–111.

[4] P.F. Isoz, J.B. Lister, P. Marmillod, A hybrid matrix multiplier for control of the TCV Tokamak, in: Proceedings of the Symposium on Fusion Technology, vol. 1264, London, 1990.

[5] O. Barana, A. Luchetta, G. Manduchi, C. Taliercio, Recent developments of the RFX control and data acquisition system, in: 4th IAEA TCM on Control, Data Acquisition, and Remote Participation for Fusion Research, San Diego, CA, 2003, in press.
[6] W. Davis, P. Roney, T. Carroll, T. Gibney, D. Mastrovito, The use of MDSplus on NSTX at PPPL 3rd IAEA TCM on Control, Data Acquisition, and Remote Participation for Fusion Research, Padova, Italy, 2001, pp. 247–252.
[7] O. Barana, A. Luchetta, G. Manduchi, C. Taliercio, Design choices for control and data acquisition in RFX, in: Proceedings of the 3rd IAEA TCM on Control, Data Acquisition and Remote Participation for Fusion Research, Padova, Italy, 2001, pp. 361–366.