

RFC Linux Transition Team
Findings and Recommendations
Report

February 3, 2003

Linux Transition Implications for RFCs

1. Executive Summary

The RFC Linux Transition Team was formed shortly after AWIPS announced it had plans to replace all the HP-UX workstations in FY 2003 and replace the data and application servers in later years (FY 2005 or later). The team was given the charge to document plans and guidelines for an efficient, effective transition to an all Linux workstation environment and to identify any issues which needed to be addressed by AWIPS for RFCs to complete the transition. The team produced one technical document with plans and guidelines for the RFCs, and this document with the issues and recommendations for AWIPS.

Based on their investigations, the RFC Linux Transition Team makes the following recommendations to the AWIPS program:

1. Allow the RFCs to implement the execution of scheduled batch operations on the Linux PX system.
2. The AWIPS Systems Engineering Team should begin working immediately with the Linux Transition Team to devise a standard solution to the issue of additional disk space for RFCs. This solution should be able to be implemented by each RFC prior to the replacement of the HP workstations.
3. A solution to the color problem caused by hosting the HP Xterminals on Linux workstations should be identified prior to the removal of the HP workstations. This solution should be acceptable to OST, OHD, and OCWWS-HSD (as proxy for the RFCs).
4. The team(s) formed to determine the solutions to providing RDBMS and GIS capabilities in an all-Linux AWIPS environment should include OHD/HL personnel to adequately consider the requirements of hydrologic software applications. The proposed solutions should be tested at an RFC to determine their effectiveness in the RFC environment.

2. Introduction

The AWIPS program has announced that the original AWIPS hardware, comprised of servers and workstations using the HP-UX operating system, will be replaced with PC equipment using RedHat Linux. The plan for this replacement includes a phased replacement of the various components over several years. The most significant components from the perspective of the hydrology software are the workstations and the DataServers (DS). Most of the interactive and user initiated applications are run on the workstations. The scheduled application executions are performed on DS which also hosts the relational database. The workstation replacement is scheduled to take place during FY03, while the DS replacement is planned for FY05.

In response to this announcement, the RFC Linux Transition Team (Appendix A) was formed to document plans and guidelines for transitioning RFC operations into the Linux operating system environment. The creation of a team was suggested at the June, 2002 Hydrology Science Conference and formed under the RFC Development Management process. The impetus for a team focused on RFCs was recognizing the operations and software at the thirteen RFCs are significantly different from those at other AWIPS sites and the transition to Linux may therefore raise issues different from those other sites. The team was also able to take advantage of the experience gained by several RFCs working with Linux on demonstration projects.

After several months of investigation and discussion the transition team produced two documents. The first is a set of technical guidelines for transitioning the data, local applications and operations of RFCs. In the second (this document), the team identifies several issues which must be resolved before RFCs are able to effectively conduct their operations in an "all Linux" environment. The main purpose of this document is to ensure these issues are recognized by all the parties involved, so solutions can be found which will allow the successful transition of RFC operations.

The issues which need to be solved are: 1) disk space for storing Linux versions of data and executables, 2) a Linux platform for the scheduled execution of NWSRFS applications, 3) the Informix RDBMS, and 4) the ArcView Geographic Information System (GIS). The first two issues are of immediate concern because they are closely tied to the workstation replacement which is expected within the next year. The last two issues are associated with software which runs on the DS, which are not scheduled to be replaced until FY05.

3. Background

This document deals mainly with transition issues for the AWIPS delivered hydrology software which runs at RFCs. This includes the Operational Forecast System (OFS), Calibration System (CS), Ensemble Streamflow Prediction (ESP), and Flash Flood Guidance System (FFGS) of the National Weather Service River Forecast System (NWSRFS), Multisensor Precipitation Estimation (MPE), and the WFO Hydrologic Forecast System (WHFS).

NWSRFS, the main RFC hydrologic forecasting software, relies upon a custom binary flat file data base used by the OFS, ESP, and FFGS components. This data base was developed and built prior to the inception of relational databases and general NWS use of Informix. Although there

are disadvantages to using a custom database, the NWSRFS has served the RFCs well for many years and does provide some advantages over proprietary relational database engines, such as: 1) it provides very fast access to data and 2) it enhances national and international technology transfer because there is no license required.

One disadvantage of the custom binary database is its hardware dependence. Binary file storage differs on PC and Workstation hardware. The primary byte, the one which denotes the type of data being stored (the big endian) for any binary data is placed at opposite ends of the data on the two platforms. So binary data written to a PC (Linux) platform cannot be read by a workstation application unless the endians are flipped. This problem is referred to as the “big endian – little endian problem”. [Note: throughout this paper, the term “Linux environment” is often used to refer to the PC hardware environment, while HP-UX is used synonymously with the Workstation architecture.]

This endian problem is easily solved if the binary data files are all of known format and if there are only a few different file formats. However, within the NWSRFS there are over 15 different file formats, and not all of the formats are fixed. The files have no known format; they are merely binary dumps of float arrays into which have been packed, floats, integers and character variables. The endians of these data cannot be arbitrarily flipped. Each data element must be interpreted and then, once the data type is known, the endians can be flipped. A two byte integer, must be treated differently than a four byte float, and characters must be treated differently again.

The Hydrology Laboratory (HL) investigated three options for how to deal with this issue in the AWIPS timeframe to port NWSRFS to Linux. More details of the three options are contained in Appendix B. HL discussed these options with members of the AWIPS Systems Engineering Team and Software Engineering Group on August 23, 2001 and then selected the Accelerated Transition solution which was implemented as part of porting all of the AWIPS delivered hydrology software to Linux.

HL decided the Accelerated Transition approach provided the quickest route to full Linux implementation and the benefits of additional computing power provided by the new PC systems. It was the most cost effective solution because it required no additional coding except that which is required by the Linux port itself. There were no transitional pieces of code which would consume substantial resources to create and have only short life spans. As a consequence of this solution, all the NWSRFS executables must be transitioned from HP-UX to Linux at the same time. Due to the lack of transitional pieces of code to identify the source and endian format of the data, the applications write binary data in the native format for the platform on which they are executing. Since many of the NWSRFS executables are executed in both scheduled and on-demand modes reading and writing the same binary data, both the scheduled and on-demand executions must be run in the same environment. (There are a few exceptions to this for operations which produce ASCII files, which are not affected by the endian problem.) Therefore, when the interactive applications begin running on Linux workstations, all of the batch operations must also be run in the Linux environment.

4. Issues

Scheduled Batch Operations

In the all HP-UX AWIPS environment, the scheduled batch operations are run on the DS using a system “cron” capability. The DS system has a failover capability which ensures if one DS computer fails, the tasks will switch and be run on the backup DS machine. This reliability is necessary to guarantee essential applications and services are available and executed when needed. Unfortunately, the phased schedule for the replacement of AWIPS equipment results in an extended period of time in which the workstations and DS will be operating with different operating systems. This presents a serious problem for the execution of scheduled NWSRFS applications because, as explained above, these must be executed in the Linux environment, which is not provided on the DS. The Linux workstations provide the correct environment, but do not have the automatic failover capability to guarantee execution of the scheduled executions. This situation is exacerbated at an RFC, which is typically not staffed at all times, so a workstation failure might go several hours without being detected, delaying manual rescheduling of the executions on a functional workstation.

OHD representatives discussed this problem with the AWIPS SET and a tentative solution for the problem was identified. The solution is to run the scheduled executions on a Linux Preprocessor system (PX) which is scheduled to be deployed at all AWIPS sites before the Linux workstations are deployed. The PX can provide both the Linux environment and the automatic failover capability required for the guaranteed execution of the scheduled NWSRFS applications.

Recommendation 1: Allow the RFCs to implement the execution of scheduled batch operations on the Linux PX system.

Disk Space

The other critical issue which needs to be addressed with the deployment of the Linux workstations is the provision of sufficient disk space to store the Linux versions of the NWSRFS executables and data. The endian issues discussed above dictate that binary files, including executables, which are written on a Linux (PC) system will be incompatible with those written on an HP-UX (Workstation) system. Thus two versions of each file will probably be necessary during the transition of operations from one environment to the other. It is important to understand either version may be stored on a disk regardless of the operating system running on the platform accessing the disk. In other words, both Linux and HP executables may be stored on either the current DS (running HP-UX) disks and new Linux workstation disks. The endian format only matters when the files are interpreted by a program which expects to find the files in a particular format.

In preparation for the transition to Linux, the LTT tried to estimate the likely amount of disk space required to support all of Linux versions of the hydrologic executables and data for the next three years. Based on the experience of the NCRFC, which is running its operations in a Linux environment and extrapolating the amounts of data currently used by other RFCs, the LTT estimated 35 GB of files (executables and data) would be required. Therefore, at the AWIPS SET meeting on October 17, 2002 OHD made a request 35 GB of currently unassigned space in

the PX disk storage be allocated for use by hydrology to deal with this issue. This request was denied because it would have required a change to the disk image at too late a point in the design schedule. Several alternatives were proposed including allocating space on the Archive Server system or using the current disk space on the DS systems to store the additional files. The LTT has investigated these alternatives and has found neither is workable. There is very little unallocated disk on the RFC Archive system and most RFCs are barely under the disk usage thresholds for the DS machines without starting to store Linux files. (The utilization percentages of /awips/hydroapps at most RFCs were recently between 85% and 90%, with a high of 94% and a low of 74%.) In addition to the eventual amount of disk space needed to support operations in a Linux environment, during the transition period, RFCs will want to operate in parallel on both the HP-UX and Linux environments to ensure consistency of the results.

In the absence of a standard solution, each RFC is likely to invent their own solution ranging from using disk space on a workstation to connecting one or more additional computers or disk drives (see CCB request AB100). This will make AWIPS configuration management and support very difficult.

Recommendation 2: The AWIPS Systems Engineering Team should begin working immediately with the Linux Transition Team to devise a standard solution to the issues of additional disk space for RFCs.. This solution should be able to be implemented by each RFC prior to the replacement of the HP workstations.

Xterminals

The current recommendation for the Linux workstations retains the HP Xterminals as the third head in the nominal AWIPS three-headed workstation. This approach will mean that the HP Xterminal will be hosted by (or boot from) its associated Linux workstation.

For a number of years, RFCs have physically separated the Xterminals from each workstation to allow additional access points to the AWIPS system. In the current configuration (HP Xterminal hosted by HP workstation), the Xterminals are capable of executing any application which can be run on a workstation. Based on experience at NCRFC with the proposed configuration, the displays of some applications on the Xterminals will be degraded. This concern has been shared with members of the AWIPS Systems Engineering Team. OST personnel have made considerable efforts to identify the cause of this problem and have identified at least one workaround which can be accomplished through modifying a software configuration file. The root cause of the problem is that the default configuration of the windows manager on the Linux workstation does not allow sufficient colors required for the hydrology related applications. The feasibility and potential side effects of implementing this workaround are being investigated.

Recommendation 3: A solution to the color problem caused by hosting the HP Xterminals on Linux workstations should be identified prior to the removal of the HP workstations. This solution should be acceptable to OST, OHD, and OCWWS-HSD (as proxy for the RFCs).

Long-term Issues

The first long-term Linux transition issue which will need to be addressed prior to replacing the HP-UX DS machines is the transition or replacement of the Informix Relational Data Base Management System (RDBMS). Many hydrologic applications depend on the hydrology database managed by this package. The LTT understands the AWIPS program plans to examine options and alternatives for Informix in a Linux environment. The two main concerns of the LTT are 1) OHD/HL personnel are included in the team determining the RDBMS solution for the Linux environment, and 2) the RDBMS solution is selected with sufficient time to allow for revision of applications and the extensive testing required to ensure all of the applications work correctly with the RDBMS solution. All hydrologic applications have been written without making use of specific Informix extensions to standards, but any differing implementation, even just a later version of the Informix package, may require some modification of the application interfaces.

The second long-term Linux transition issue deals with the replacement of the ArcView Geographic Information System (GIS). This package is part of the AWIPS baseline currently used by a few applications, but several additional major applications which require GIS data and functionality are currently under prototype development and testing. The vendor of ArcView, ESRI, has announced it has no plans to offer this product in the Linux environment, so a replacement needs to be selected or another avenue for use of ESRI software needs to be included in AWIPS. As with the RDBMS, the LTT has the same two concerns: 1) OHD/HL personnel need to be involved in the selection process and 2) the selection must be made in a timely manner.

Recommendation 4: The team(s) formed to determine the solutions to providing RDBMS and GIS capabilities in an all-Linux AWIPS environment should include OHD/HL personnel to adequately consider the requirements of hydrologic software applications. The proposed solutions should be tested at an RFC to determine their effectiveness in the RFC environment.

Appendix A: Team members

Steve Buan	Team Leader	NCRFC
John Halquist		NCRFC
Billy Olsen		ABRFC
Bill Lawrence		ABRFC
Adrian Faimon		MBRFC
John Kuhn		OCWWS-HSD
Joe Gofus		OHD

Appendix B: Details of Proposed Solutions to Endian problem for NWSRFS

The Hydrology Lab (HL) identified three possible solutions with respect to the current Linux port. The first is to port the NWSRFS and maintain two databases which are kept synchronised. The second is to port the NWSRFS to Linux in such a way that it will flip the endians on the fly. The third is to port the NWSRFS to Linux and provide the RFCs with an accelerated transition to Linux.

Dual Database Solution

A Dual Database Solution requires that the NWSRFS be ported to Linux and that a new process be developed to keep a Linux database and an HP workstation database synchronized. This solution becomes difficult because the NWSRFS databases are accessed and updated by multiple users simultaneously. In addition, the order of the updates is critical to successful River Forecasting – model runs must be made in an upstream to downstream order. With multiple forecasters updating different portions of large basins the management of the upstream downstream computational order becomes complex. Substantial design and testing would be required to develop an effective synchronization process.

Flip on the Fly Solution

The Flip on the Fly solution requires that all the read and write routines used in the NWSRFS be updated to a) determine on which operating system they are running b) flip the endians as required. Because of the unformatted file structure, the on the fly flip would have to take place at a relatively high level – where the records are decoded. Because the flipping would have to take place where the records are understood the number of places where the flipping would be required is large (and unknown at this time). The flipping could not be limited to the few routines where the files are actually read and written. Substantial testing would be required to insure that no database corruption would occur when the flipped records were written back to the database.

Accelerated Transition Solution

The Accelerated Transition solution requires that the NWSRFS be ported to Linux and that the HL provide rapid support for the accelerated transition to Linux. This solution requires shortening the transition period over which the RFCs will move to the Linux operating system and the support of the AWIPS program to succeed.