

Supercomputing Center Management using AIRS

Robert A. Ballance, Jared Galbraith, and Roy Heimbach

High-Performance Computing, Education and Research Center
The University of New Mexico
Albuquerque, New Mexico

Introduction

Running a large University supercomputing center teaches many lessons, including the need to centralize data collection and analysis, to automate system administration functions, and to enable users to manage their own projects. AIRS evolved in response to these pressures.

The setting

The HPCERC at the University of New Mexico, is a leading academic site for high-performance computing and scientific programming. When Los Lobos, the Center's Linux-based 512-processor supercluster, came on-line in June, 2000, HPCERC ranked among the largest five academic supercomputing centers in the U.S. The Center also operates a 128 processor production Linux supercluster on behalf of NCSA.

Locally, the Center maintains three 32-processor research superclusters from IBM and VA Linux Systems, Silicon Graphics visualization servers, a novel visualization cluster for parallel graphics, a visualization laboratory, and a cluster of workstations, all in support of a broad range of scientific processing. The Center owns and operates two Access Grid studios `ag.ballance00:_nation_comput_scienc_allian_acces_grid` for use in Internet-based teleconferencing research. All of our production and research systems run either Linux, BSD, AIX, or IRIX, with Linux the dominant OS. Staff members choose from a menu of Linux, BSD, Mac, or Windows as their primary workstation OS.

Federally-funded research scientists from across the U.S. use HPCERC resources as part of their research programs. Within UNM, HPCERC assists over 20 associated faculty and their students, from the Colleges of Arts and Sciences, Engineering, Fine Arts, and the School of Medicine in their research and teaching programs. Nationally, HPCERC partners with other Alliance member institutions in developing, deploying, and promoting high-performance computing technologies.

The Problem

Imagine this, then:

2000 Robert A. Ballance, Jared Galbraith, and Roy Heimbach

- Multiple clusters, each with their own accounting systems;
- Distinct and separate naming domains (`alliance.unm.edu` and `ahpcc.unm.edu`) within a single facility;
- Required daily usage reporting via electronic data transfer to NCSA
- Dozens of users and projects, many administered remotely;
- Active and historical data scattered in multiple locations, including flat files, desktop databases, and on paper.

The Solution: Take a deep breath of ... AIRS

What's a manager to do? Reach for the open source and get to work. Two years later, the result is AIRS, an open source solution for managing today's supercomputing center. The name was originally just the "Integrated Reporting System" --- chosen for its acronym and its ability to insinuate itself into virtually every aspect of Center management.

Core features of the system, including daily reporting of usage information, have been online since early 2001. Other features are in various stages of deployment. We are working to simplify installation and deployment at other sites. Distribution tarballs will be available at [\HREF{airshq.hpcerc.unm.edu/}](http://airshq.hpcerc.unm.edu/).

This article provides a brief overview of AIRS: its capabilities, design, and evolution.

Design Pressures

The development of AIRS began several months after our first large Linux cluster (RoadRunner) came online as a production computing environment. The RoadRunner cluster was funded jointly by NCSA and UNM. As a production machine, we needed to be able to account for the usage of the machine. For example, daily usage had to be reported to NCSA. Not all usage, but just the usage for those projects with specific allocations. However, as soon as we could gather information for our sponsors, we immediately began to use it ourselves. It is a comfort to receive mail each day summarizing the usage for the day before. It was a small leap to add daily, weekly, biweekly, monthly, and quarterly reports.