DEVELOPING & MANAGING A LARGE LINUX FARM - THE BROOKHAVEN EXPERIENCE

A. Chan*, R. Hogue, C. Hollowell, S. Misawa, O. Rind, J. Smith, A. Withers, T. Wlodek
Brookhaven National Laboratory, Upton, NY 11973, USA

Abstract

This presentation describes the experience of the RHIC/ATLAS Computing Facility (RCF/ACF) in building and managing its 2,700+ CPU (and growing) Linux Farm at Brookhaven National Laboratory over the past 6+ years. We describe how hardware cost & configuration, resource requirements, software support and other considerations have played a role in the process of steering the growth of the RCF/ACF Linux Farm in the midst of an ongoing transition from being a generally local resource to a global, Grid-aware resource within a larger, distributed computing environment.

INTRODUCTION

The RHIC Computing Facility (RCF) is a large scale data processing facility at Brookhaven National Laboratory (BNL) for the Relativistic Heavy Ion Collider (RHIC), a collider dedicated to high-energy nuclear physics experiments. It provides for the computational needs of the RHIC experiments, including batch, mail, printing and data storage. As the U.S. Tier 1 Center for ATLAS computing, the ATLAS Computing Facility (ACF) provides for the computational needs of the U.S. collaborators in ATLAS. Increasing support for Grid-like technologies is transforming the RCF/ACF from a local into a globally available computing resource, with growing design and operational complexity that requires increasing staffing levels (see Fig. **??**) to handle additional responsibilities.

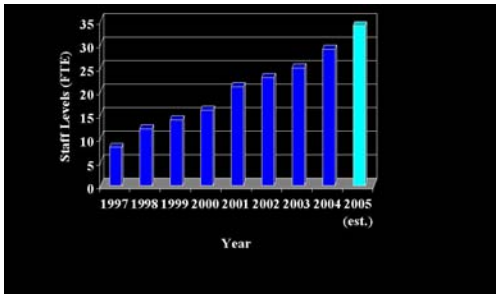

Figure 1: Staffing Levels at the RCF/ACF.

* awchan@bnl.gov

THE LINUX FARM – PAST, FUTURE AND PRESENT

The Linux Farm is the main source of CPU and distributed storage resources and the user interface to other services such as HPSS (mass storage), NFS (central storage) and AFS within the RCF/ACF. The fast growth of the Linux Farm cluster and the appearance of Grid-like technologies have a direct impact on the operation, management and availability of the Linux Farm. We discuss below how these changes are creating changes in management philosophy, hardware requirements, software support and security-related matters.

Change in Management Philosophy

The growth of the Linux Farm cluster makes it imperative to have a scalable, automated monitoring (see Fig. **??**) and management system with capabilities such as remote power management and predictive hardware failure analysis. The ability to quickly and accurately determine the status of the cluster is important for prompt problem resolution and service restoration. Servers are no longer treated as individual servers but as part of a larger collective of servers. High availability is based on the sheer number of identical servers, not on the 24-hour availability of individual servers. Increasingly larger clusters are only manageable if servers are identically built, hardware and software-wise. For this reason, building specialized, custom servers within the Linux Farm is discouraged.

Evolution in Hardware Requirements

Early versions of the Linux operating system (OS) coincided with the appearance of desktops and servers built with commodity hardware in the early 90's, and rack-mounted servers generally appeared in computer centers in the mid 90's. The growth of the Linux Farm also coincided with a drop in server price (see Fig. **??** and **??**). A parallel drop in price for server-based disk storage (see Fig. **??** and **??**) and the Grid-driven interest for distributed storage has changed the focus of our hardware acquisitions from being solely CPU-driven to CPU & storage-capacity driven, since, based on our experience, centralized fiber-channel storage is several times more expensive than server-based storage (SCSI, SATA or IDE technologies). Fig. **??** shows how we increased the storage-capacity per server in our most recent hardware acquisition cycles. To keep hardware reliability high, we have restricted our acquisitions to high-quality vendors which can deliver servers and scalable

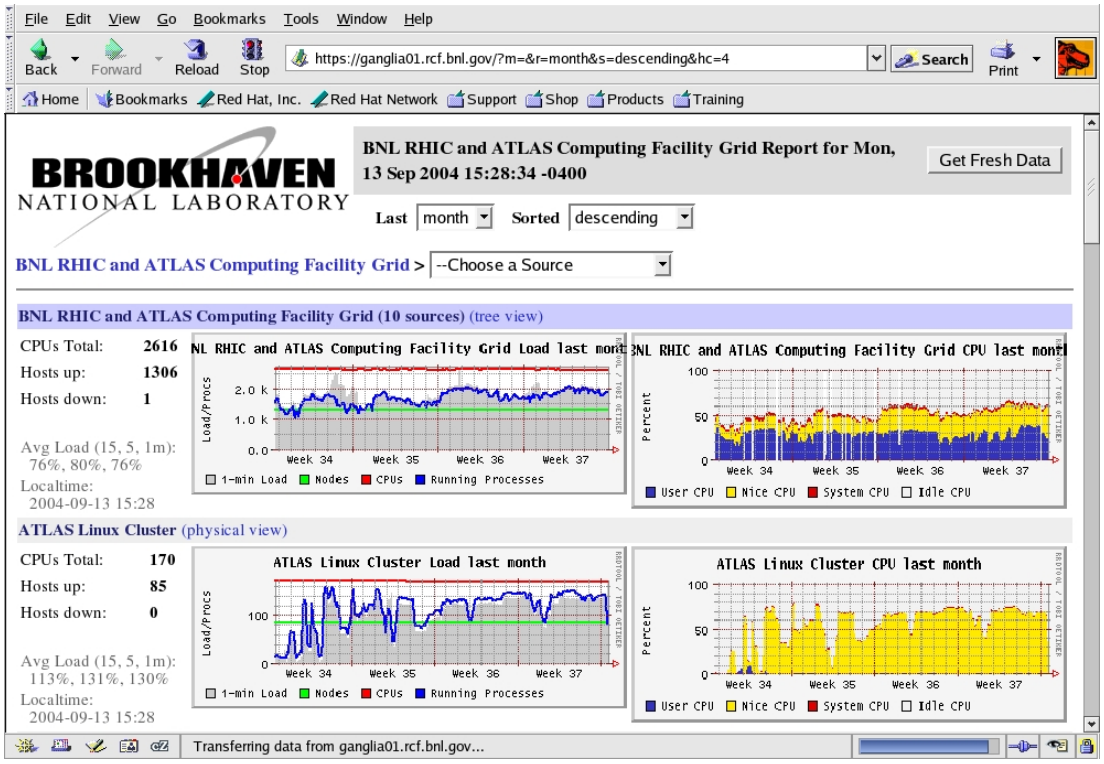

Figure 2: The ganglia [**?**] monitoring software at the RCF/ACF.

cluster management software with minimal labor-intensive maintenance requirements.

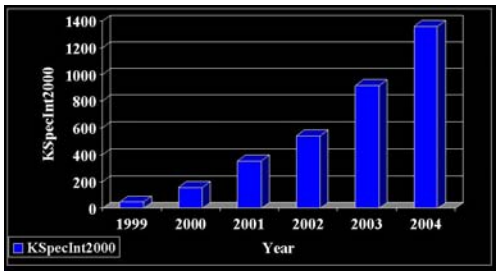

Figure 3: Growth of the Linux Farm over time.

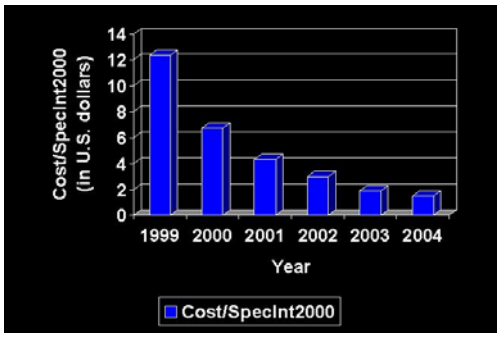

Figure 4: Server cost based on price/performance.

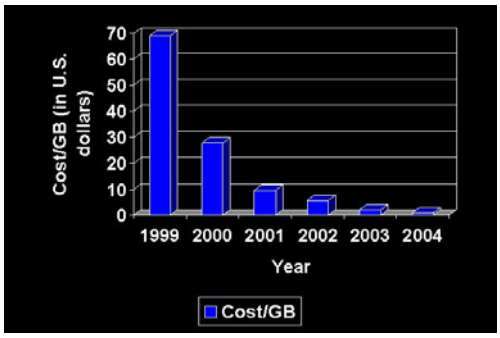

Figure 5: Cost of local storage over time.

Evolution in Software Support

Gradual changes in software support at the Linux Farm are also occurring to keep up with changing needs. In the recent past, locally-built, commercially available or open-source software was chosen on the basis of familiarity, local requirements and ease of use. The fast growth of the Linux Farm and the transformation to a globally available resource make software scalability and inter-operability with other sites more critical than in the past, and it is leading the Linux Farm towards a loose set of basic, mutually acceptable software packages that are common among different sites. Criteria for defining this common set of software packages is still being defined, but open-source, wide acceptance and scalability are important considerations at the RCF/ACF. Table **??** summarizes some of the major

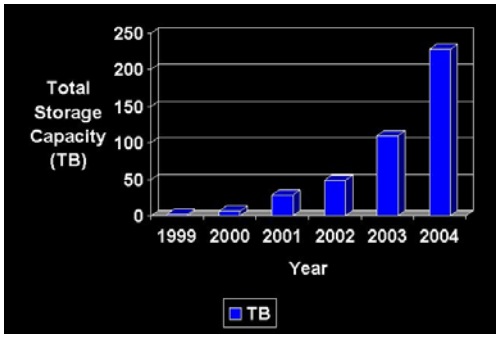

Figure 6: Growth of local storage capacity in the Linux Farm cluster.

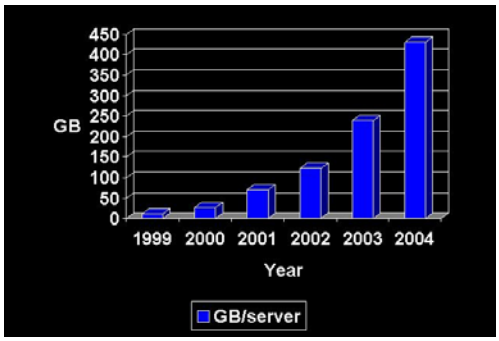

Figure 7: Growth of storage capacity per server in the Linux Farm.

transitions in software package support within the Linux Farm over the last few years.

Table 1: Software Evolution

Package	Old	New	Date
OS	RH Linux 8	SL 3	2004
Batch	Home-Built	Condor/LSF	2004/2000
Monitoring	Home-Built	Ganglia	2003
Security	NIS	K5/GSI	2003/2004
Dist. Stor.	–	HRM/dCache	2004/?

An example of a new software package that meets our evolving requirements is the Condor [**?**] batch system. Fig. **??** shows the monitoring interface for Condor that allows the RCF/ACF track usage of this batch system.

Security

As in many computer facilities, the RCF/ACF started with telnet/NIS in the 90's. Security concerns stemming from frequent cyber-attacks at U.S. and foreign computer centers (many of linked to scientific laboratories) caused the RCF/ACF to devote more resources to cyber-security. It led to the installation of firewalls, creation of gatekeeper servers and a migration to **ssh** (secure shell) as the access method at the RCF/ACF. The Linux Farm is currently only accessible through these dedicated gatekeepers. Security updates for the Linux Farm OS have become more common-place, and OS upgrades in general must pass higher security standards than in the past.

There is on-going transition to Kerberos 5 (K5) authentication, with the eventual phasing out of NIS authentication. Testing and support of GSI authentication are also growing, mainly among those services where this form of authentication is required.

Currently, authorization and authentication are controlled by the local site. A migration to GSI will require a central CA (Certificate Authority) and regional VO's (Virtual Organization) for authorization and initial authentication. The RCF/ACF also foresees the possibility of accepting certificates from a finite number of CA's (rather than a single one). The local sites would perform final authentication before granting access to local resources. The transition from complete to partial control over security issues has been and will continue to be difficult due to concerns about loss of control over security issues at local sites.

SUMMARY

The RCF/ACF Linux Farm is undergoing many changes in its transition from a local to a global facility.

Building reliable computer clusters with commodity hardware with open-source software has become prevalent and affordable. Major computer vendors have been offering an increasing variety of servers built with commodity hardware, insuring that prices will stay competitive for the foreseeable future. Distributed storage has also become increasingly feasible as cost for disk storage drops, although major questions about management software remain unanswered. Tests with **dCache** [**?**] look promising. Inter-operability with remote sites has and will continue to play an increasingly important role in our choice of software packages.

A transition to a globally available computing resource will likely take longer and be more difficult than generally expected. Changes in hardware and software requirements must be complemented by changes in management and operational philosophies, particularly with regards to security and user access issues.

ACKNOWLEDGEMENTS

The RCF/ACF would like to thank BNL's Physics Department, Information Technology Division (ITD) and the United States Department of Energy (DOE) for their support.

REFERENCES

[1] http://ganglia.sourceforge.net/

[2] http://www.cs.wisc.edu/condor

[3] http://www.dcache.org

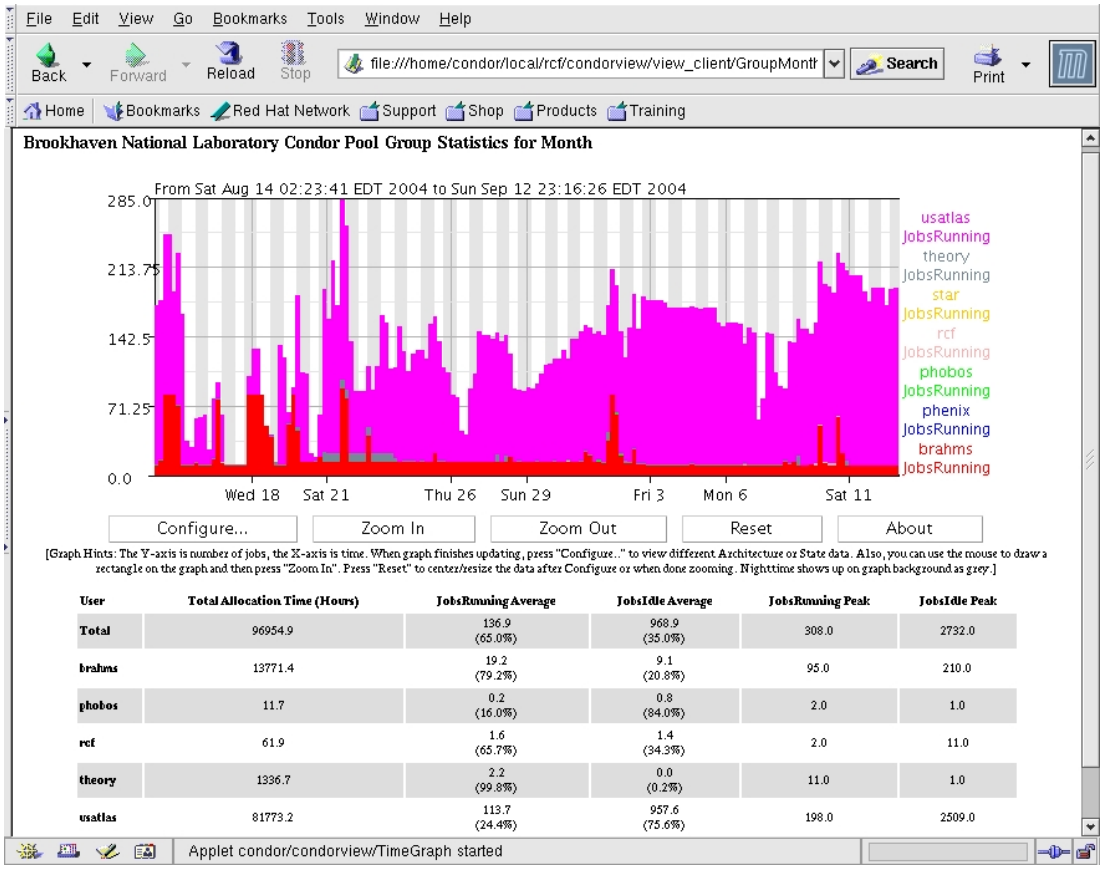

User	Total Allocation Time (Hours)	JobsRunning Average	JobsIdle Average	JobsRunning Peak	JobsIdle Peak
Total	96954.9	136.9 (65.0%)	968.9 (35.0%)	308.0	2732.0
brahms	13771.4	19.2 (79.2%)	9.1 (20.8%)	95.0	210.0
phobos	11.7	0.2 (16.0%)	0.8 (84.0%)	2.0	1.0
rcf	61.9	1.6 (65.7%)	1.4 (34.3%)	2.0	11.0
theory	1336.7	2.2 (99.8%)	0.0 (0.2%)	11.0	1.0
usatlas	81773.2	113.7 (24.4%)	957.6 (75.6%)	198.0	2509.0

Figure 8: The graphical monitoring interface for the condor batch system.